

Smarter Geothermal Field Development with Agentic AI System, GAIA

Randy Harsuko^{1,2}, Zhengfa Bi², Nori Nakata²

¹King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

²Lawrence Berkeley National Laboratory, Berkeley, CA

mochammad.randycaesario@kaust.edu.sa

Keywords: AI agents, geothermal, retrieval augmented generation

ABSTRACT

Geothermal field development typically involves complex processes that require multi-disciplinary expertise in each process. Thus, decision-making often demands the integration of geological, geophysical, reservoir engineering, and operational data under tight time constraints. We present Geothermal Analytics and Intelligent Agent, or GAIA, an AI-based system for automation and assistance in geothermal field development. GAIA consists of three core components: GAIA Agent, GAIA Chat, and GAIA Digital Twin, or DT, which together constitute an agentic retrieval-augmented generation (RAG) workflow. Specifically, GAIA Agent, powered by a pre-trained large language model (LLM), designs and manages task pipelines by autonomously querying knowledge bases and orchestrating multi-step analyses. GAIA DT encapsulates classical and surrogate physics models, which, combined with built-in domain-specific subroutines and visualization tools, enable predictive modeling of geothermal systems. Lastly, GAIA Chat serves as a web-based interface for users, featuring a ChatGPT-like layout with additional functionalities such as interactive visualizations, parameter controls, and in-context document retrieval. To ensure GAIA's specialized capability for handling complex geothermal-related tasks, we curate a benchmark test set comprising various geothermal-related use cases, and we rigorously and continuously evaluate the system's performance. We envision GAIA as a pioneering step toward intelligent geothermal field development, capable of assisting human experts in decision-making, accelerating project workflows, and ultimately enabling automation of the development process.

1. INTRODUCTION

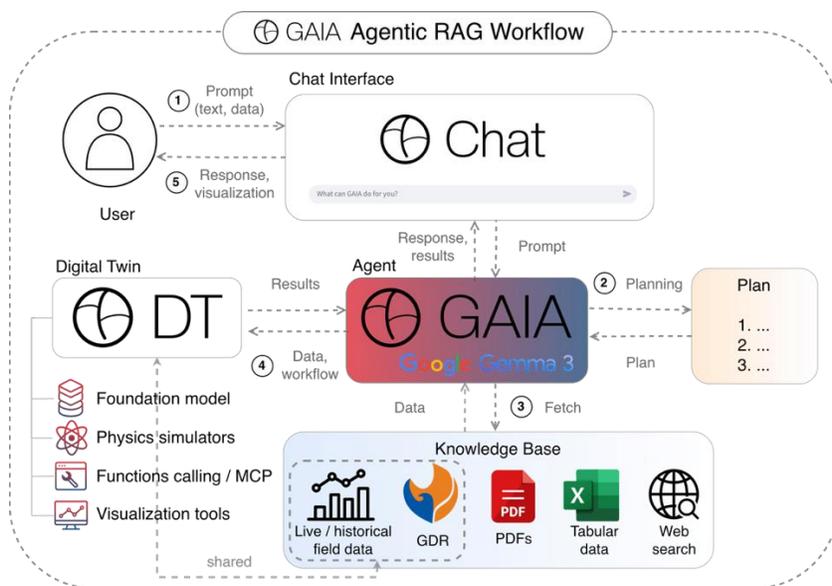


Figure 1: GAIA Agentic RAG Workflow. Users interact with GAIA Chat interface by supplying prompts and supplementary data. The GAIA Main Agent will first think through the prompt and make a step-by-step plan to solve the given problem. If necessary, the system will fetch data from its knowledge base consisting of PDFs of papers/proceedings, tabular data, and/or external data sources (via web search and/or GDR). The collected data and the formulated workflow will be forwarded to GAIA DT for the actual processing through its set of tools. Finally, the Main Agent will formulate the final response based on the gathered information and multi-step analyses.

Geothermal energy is a renewable energy source that harnesses heat from the Earth's interior for electricity generation and direct use applications. It has the potential to provide a significant portion of the world's energy needs while reducing greenhouse gas emissions. However, geothermal field development is a complex process that involves multiple stages, including resource assessment, exploration, drilling, production, and reservoir management and optimization. Each stage requires specialized knowledge and expertise in various

fields such as geology, geophysics, reservoir engineering, and environmental science (Gensheng et al., 2024). The complexity of geothermal field development presents challenges for decision-making, as it requires experts to integrate diverse data sources, interpret complex geological and geophysical information, and make informed choices under time constraints. To address these challenges, there is a growing interest in leveraging artificial intelligence (AI) and machine learning (ML) techniques to automate and assist in geothermal field development (e.g., He et al., 2020; Nakata et al., 2025; Gutiérrez-Oribio et al., 2025; Montes et al., 2025). Recently, AI agents have shown promise in various domains, including natural language processing (NLP), computer vision (CV), and robotics (Chan et al., 2024).

AI agents, which are sometimes also referred to as "Deep Research Agents", are LLM-powered systems that integrate dynamic reasoning, adaptive planning, multi-iteration external data retrieval and tool use, and comprehensive analytical report generation for research purposes (Huang et al., 2025b). They are meant to assist users in automating complex tasks, such as data analysis, literature review, and experimental design. Some examples of recent AI agents include OpenAI Deep Research (OpenAI), Alita (Qiu et al., 2025), Pangu DeepDiver (Shi et al., 2025), etc. For a more comprehensive review of AI agents, please refer to Huang et al. (2025b).

In the geoscience sphere, several LLM-based systems have been proposed. ChatClimate (Vaghefi et al., 2023) is a web-based chatbot app built for querying climate-related information based on the latest IPCC report. It leverages the standard retrieval-augmented generation (RAG) pipeline to reduce hallucinations and provides more accurate and factual answers on climate. Pioneering LLM-based geological map analysis, Huang et al. (2025c) promote PEACE, an agentic system with GPT-4o as a base model that excels in various geological map-understanding tasks. They also curate a benchmarking dataset for evaluating AI agents in geologic map understanding. AskGDR (Weers et al., 2024) is the first LLM-based app introduced to the geothermal community. The chatbot-styled web app was trained on metadata from collections of geothermal datasets in the Geothermal Data Repository (GDR) and was designed to assist researchers in obtaining geothermal datasets that suit their needs. For seismic exploration applications, Kanfar et al. (2025) is the first to introduce an AI agent that specializes in building seismic processing workflows. Rather than manually writing Madagascar (Fomel et al., 2013) scripts, the agent can automatically generate and execute the scripts based on user prompts, significantly reducing the time and effort required for seismic data processing (Kanfar et al., 2025). However, currently, there is still limited, if not any, research on agentic AI systems that handle a complete, end-to-end geothermal field development workflow.

The objective of this work is to develop an AI-based system, GAIA – Geothermal Analytics and Intelligent Agent, for automation and assistance in geothermal field development. GAIA addresses the complexity and challenges of decision-making, which often requires integrating geological, geophysical, reservoir engineering, and operational data under time constraints. The system is designed to assist experts throughout the workflow, from data analysis and simulation to decision support and project automation. To the best of our knowledge, GAIA is the pioneering effort in building an agentic AI system for geothermal project assistance.

2. METHODS

GAIA comprises three components: GAIA Agent, GAIA Chat, and GAIA Digital Twin or DT, forming an agentic RAG workflow. GAIA Agent, powered by a pre-trained LLM, autonomously queries knowledge bases, executes domain-specific subroutines, and orchestrates multi-step analyses. GAIA DT integrates classical and surrogate physics models with interactive visualization tools, while GAIA Chat provides a web-based interface with features such as interactive visualizations, parameter controls, and context-aware document retrieval. Figure 1 illustrates the GAIA Agentic RAG workflow. Each of these components will be discussed in detail next.

2.1 GAIA Agent

The core component of agentic AI systems is an AI agent, which is based on a pre-trained LLM. In the GAIA framework, we define a main agent that is responsible for planning, orchestrating, and executing tasks. Therefore, the minimum requirement of the pre-trained LLM is that it needs to understand instructions (e.g., via instruction tuning (Peng et al., 2023)) and generate human-like text. In this work, we use the open-weight model of Gemma 3 from Google Deepmind (Team et al., 2025), which is a 27B parameters LLM model that was trained on 14T tokens of text and image data. We specifically use the Gemma3-27B-IT (instruction-tuned) variant of the model, which is designed to follow user instructions and generate coherent and contextually relevant responses. We acknowledge that other LLMs can be used as the main agent, such as OpenAI's GPT-5, Meta's Llama 4, and Anthropic's Claude 4. We choose Gemma 3 for its open-weight availability and its strong performance in instruction-following tasks, comparable to one of the earlier flagship LLMs, Gemini 1.5 pro (Team et al., 2024). Moreover, smaller LLMs can easily fit into consumer-grade GPUs, making it more accessible for users to run GAIA on their local machines and edge devices in case there is no internet connection (such as in remote geothermal fields).

The main agent operates in a combination of the Chain-of-Thought (CoT, Wei et al. (2022)) and Reaction+Act (ReAct, Yao et al. (2023)) paradigms. Specifically, the main agent will first think through the prompt and make a step-by-step plan to solve the given problem. In each step, the main agent will decide whether to deduce from the knowledge base or execute tools via the GAIA DT. Then, for every step, the main agent will analyze whether the step has been completed and the objective of the step has been achieved. These processes are done iteratively until the main objective, which is to answer the user's prompt, is achieved. The main agent will then formulate the final response based on the gathered information and the multi-step analyses.

If necessary, GAIA Agent will fetch data from its knowledge base consisting of PDFs of papers/proceedings, tabular data, and/or external data sources (via web search and/or GDR). We leverage the RAG paradigm (Lewis et al., 2020) to enhance the agent's ability and efficiency in retrieving relevant information from the knowledge base. The key to the RAG pipeline is an embedding model, which is responsible for converting multi-modal data (text, images, etc.) into vector representations. A strong embedding model should be capable of capturing the semantic meaning of the data and generating high-quality embeddings that can be used for similarity search. In this work, we use the open-weight model of Jina v4 (Günther et al., 2025), the latest open-weight embedding model with multi-modal (text and

images) support. The embedding model is used to convert the multi-modal data into vector representations, which are then stored in a vector database for efficient retrieval.

GAIA RAG system Originally, RAG systems were designed to handle text data only (Lewis et al., 2020). First, texts are tokenized and converted into vector representations using an embedding model. The vector representations are then stored in a vector database, which allows for efficient similarity search and retrieval. When a user submits a query, the query is also converted into a vector representation using the same embedding model. The vector database is then queried to find the most similar texts to the query based on their vector representations using a similarity metric (e.g., cosine similarity). The retrieved texts are then used as context for the LLM to generate a response. However, in geothermal field development, data comes in various formats, including text documents (research papers, reports, etc.), images (figures, maps, etc.), and tabular data (CSV files, spreadsheets, etc.). Therefore, we extend the traditional RAG system to handle multi-modal data by leveraging a multi-modal embedding model. For an efficient retrieval process, we store all the vector representations in a single vector database, along with their metadata to distinguish between different data types.

For the internal knowledge base, we have collected over 5,000 papers and proceedings related to geothermal field development. Prior to embedding, we preprocess the documents by splitting them into smaller chunks of text (equivalent to 2,048 tokens) using Docling (Livathinos et al., 2025) to ensure that the embedding model can effectively capture the semantic meaning of each chunk. The chunked texts are then converted into vector representations using the embedding model and stored in a vector database (we use LanceDB, <https://lancedb.github.io/lancedb/>). For images, we use the same embedding model to convert the images into vector representations, which are also stored in the vector database. For tabular data, we treat them as images and convert them into vector representations using the embedding model. Whenever the main agent needs to retrieve information from the knowledge base, it will query the vector database using the embedding model to find the most relevant chunks of text and images. The retrieved chunks will then be used as context for the main agent to generate a response. In the future, we plan to extend the knowledge base to include more data sources, such as web search results and GDR data, to further enhance the agent's ability to retrieve relevant information.

2.2 GAIA DT

GAIA DT is regarded as a digital twin in a broad sense, encompassing not only physics-based simulations but also data processing pipelines and AI-assisted workflows that connect modeling to real-world field operations and multi-modal diverse datasets. It integrates classical and surrogate physics models with domain-specific subroutines and visualization tools to enable predictive modeling of geothermal systems. Designed as a modular and extensible platform, GAIA DT allows users to incorporate new models, data streams, and subroutines as needed. The current version of GAIA DT includes the following components:

- **Seismic waveform analysis:** GAIA DT includes a set of tools for seismic waveform analysis, which can be used for tasks such as phase picking, event location estimation, magnitude estimation, and seismicity forecasting. These tools are mainly based on the ObsPy (Beyreuther et al., 2010) library, which is a widely used open-source library for seismic data analysis. The module can handle various seismic data formats, such as SAC, SEED, and MiniSEED. We plan to incorporate more advanced processing tools in next versions of GAIA DT, such as ML-based tools (Nakata et al., 2025) and real-time processor (Nakata et al., 2024).
- **Visualization tools:** Since there are many visualization types in geothermal field development, we let an AI agent that specializes in producing visualizations handle these tasks. Specifically, we use a variant of the Gemma 3 model that is fine-tuned on coding tasks, called CodeGemma. We then prompt the CodeGemma model to generate Python code for visualizations based on the main agent's (Section 2.1) requirements. The generated code will then be executed internally to produce the desired visualizations. The visualizations are based on Plotly, a popular Python library for creating interactive plots and dashboards.

Although we realize that many other components need to be added to GAIA DT, we first choose to focus on seismic waveform analysis and visualization tools in the current version of GAIA. This is because seismic monitoring is a crucial aspect of geothermal field development, as it provides valuable information about the subsurface structure and dynamics of geothermal reservoirs. Moreover, seismic data is often complex and requires specialized knowledge and expertise to analyze and interpret. In the future, we plan to extend GAIA DT to include more components, such as geophysical inversion, reservoir simulation models, risk assessment, and economic analysis. Nevertheless, the key feature of GAIA DT is the autonomous selection of the provided tools and their parameters within, which are both tedious and critical tasks, especially in a fast-paced environment of geothermal monitoring. Through a correct prompting, users could also manually interfere and modify the parameters of the selected tools at their convenience.

2.3 GAIA Chat

User interface (UI) is an important component in AI agent systems, as it serves as the primary point of interaction between users and the AI agent. A well-designed UI can enhance user experience, facilitate effective communication, and enable users to leverage the capabilities of the AI agent more efficiently. Therefore, we developed GAIA Chat, a web-based interface that allows users to interact with GAIA Agent and GAIA DT seamlessly.

For ease of testing, prototyping, and deployment, we utilize Streamlit (streamlit.io) to build a web-based app for GAIA. Streamlit is a Python library that simplifies the creation of interactive web applications, particularly for data science and machine learning projects. It allows users to transform Python scripts into shareable web apps with minimal effort and without requiring front-end web development experience (HTML, CSS, JavaScript) and contains various templates and built-in widgets. For more information on Streamlit, please refer to Khorasani et al. (2022) or their website (streamlit.io). Note that the GAIA Chat interface can be ported to production-grade frameworks like JavaScript in the future for more flexible customization, but the concept of the features will be the same, as explained below.

We adopt a layout similar to the ChatGPT web/desktop app (Figure 2). The main window of the GAIA web interface features components of a standard chatbot, like a chatbox to type and send messages, and a conversation window to display message history between a user and GAIA. Additionally, to support thorough geothermal analytics, we added the following features:

- File upload button: supports any file extensions to account for a diverse type of geothermal data. However, currently, we focus on support for structured tabular data (like CSV). Note that users could also upload files via drag-and-drop to the chatbox.
- Search setting panel: to control the number of returned search items from the knowledge base. We allow users to separately choose the maximum number of documents and images for better flexibility.
- Interactive figure output: most figures displayed by GAIA will be interactive. Thanks to Plotly, a widget-focused Python plotting library, users can pan, zoom, and hover over the generated figures. Users can also export the figures to PNG or PostScript (PS) according to their needs.

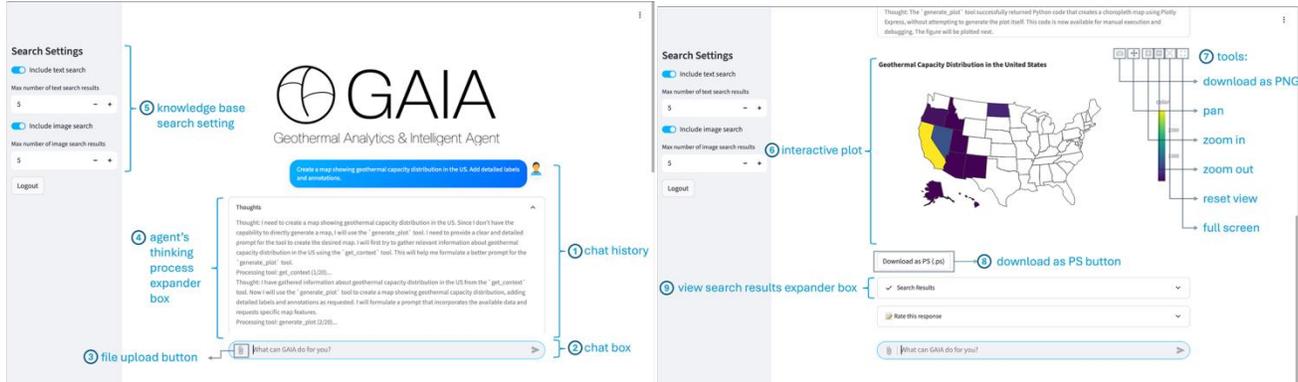


Figure 2: GAIA Chat interface. The main window features a chatbox to type and send messages, a conversation window to display message history, a file upload button, and a search setting panel to control the number of returned search items from the knowledge base. The interactive figure output allows users to pan, zoom, and hover over the generated figures, and export them to PNG or PostScript (PS).

3. EXPERIMENTS

In this section, we will discuss several use cases of GAIA in geothermal field development. These use cases are designed to demonstrate the capabilities of GAIA in automating and assisting various tasks in geothermal project workflows. The use cases include seismic waveform phase picking, event location estimation, event magnitude estimation, seismicity forecasting, and end-to-end events catalog updating and forecasting.

3.1 Use case 1: Seismic waveform phase picking

Seismic event detection and phase picking are key components in geothermal seismicity monitoring. Conventionally, amplitude-based methods dominate the seismic event detection process, which is often followed by manual phase picking. However, these methods are often time-consuming and labor-intensive, especially in geothermal fields with high seismicity rates. GAIA can automate this process by integrating seismic waveform data analysis and phase picking algorithms into its workflow.

There are two phase picking methods currently implemented in GAIA:

- The STA/LTA (Short-Term Average/Long-Term Average) method, which is a widely used trigger-based method for seismic event detection (Allen, 1978). It calculates the ratio of short-term to long-term average amplitudes of seismic waveforms to identify potential seismic events, and;
- The machine learning-based phase picking method, which utilizes a pre-trained model to classify seismic waveforms into different phases (P-wave and S-wave, see Nakata et al. (2025)). This method can be more accurate and efficient than traditional STA/LTA methods, especially in complex seismic environments. Moreover, unlike STA/LTA, this method does not rely on input parameters, such as the short-term and long-term window lengths, which can be difficult to determine in practice.

The GAIA Agent can autonomously decide which method to use based on the characteristics of the seismic data and the user's preferences provided by GAIA Chat. The default picking function is the STA/LTA method, as it is more widely used and easier to interpret. The user can also choose to override the default method by specifying the desired phase picking method in the prompt. Figure 3 shows an example of GAIA's phase picking workflow, where the user uploads a seismic waveform file and requests phase picking. GAIA Agent will then analyze the waveform data, apply the selected phase picking method, and return the results in an interactive plot.

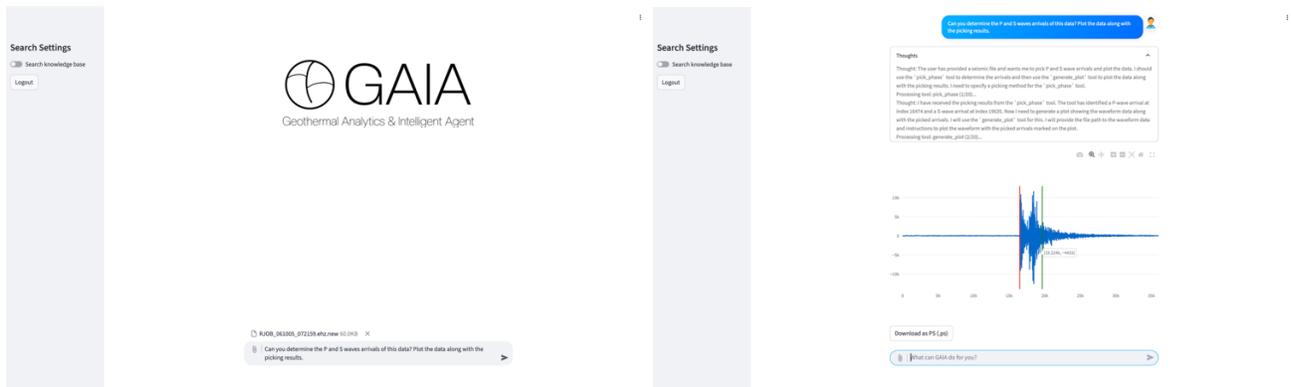


Figure 3: Example of GAIA's seismic waveform phase picking workflow. The user uploads a seismic waveform file and requests phase picking (a). GAIA Agent will then analyze the waveform data, apply the selected phase picking method, and return the results in an interactive plot (b).

3.2 Use case 2: Event location estimation

Event location estimation is a crucial step in seismic monitoring, as it determines the spatial coordinates of seismic events based on the detected phases. GAIA can automate this process by integrating various location estimation algorithms, such as the double difference method (HypoDD, Waldhauser (2001)) and the non-linear location method (NonLinLoc, Lomax et al. (2000, 2014)). These methods utilize the detected phases and their arrival times to calculate the event's location in three-dimensional space.

For simplicity, we currently implement the Geiger method (Geiger, 1912) as the default location estimation method in GAIA. The Geiger method is a linear location method that uses the arrival times of P- and S-waves to estimate the event's location. It is widely used in seismic monitoring due to its simplicity and effectiveness. However, in the future, we will add more options (like HypoDD and NonLinLoc), and users can choose other location estimation methods by specifying the desired method in the prompt.

Figure 4 shows an example of GAIA's event location estimation workflow, where the user uploads a seismic waveform file and requests event location estimation. GAIA Agent will then analyze the waveform data, apply the selected location estimation method, and return the location (latitude, longitude, and depth) and the origin time of the event.

3.3 Use case 3: Event magnitude estimation

Event magnitude estimation is another important aspect of seismic monitoring, as it quantifies the energy released by a seismic event. GAIA can automate this process by integrating various magnitude estimation algorithms, such as the Richter scale (Richter, 1935) and the Moment Magnitude (M_w) scale (Hanks and Kanamori, 1979). These methods utilize the detected phases and their amplitudes to calculate the event's magnitude.

In the current version of GAIA, we implement the Richter scale, also known as the local magnitude (ML), as the default magnitude estimation method. The Richter scale is a logarithmic scale that measures the amplitude of seismic waves and is widely used in seismic monitoring. However, it has limitations in accurately estimating the magnitude of large events and events with complex waveforms. Moreover, the local magnitude parameters, such as the distance from the station to the event and the station correction factor, are often difficult to determine in practice. Therefore, we will add more options (like Moment Magnitude), and users can choose other magnitude estimation methods by specifying the desired method in the prompt.

Figure 5 shows an example of GAIA's event magnitude estimation workflow, where the user uploads a seismic waveform file and requests magnitude estimation. Since the magnitude estimation requires the event's location, GAIA Agent will first estimate the location (Section 3.2). Then, GAIA will apply the selected magnitude estimation method. The final response will be the event's magnitude.

3.4 Use case 4: Seismicity forecasting

Seismicity forecasting is a challenging task that aims to predict the occurrence of future seismic events based on historical data. GAIA can assist in this process by integrating various forecasting models, such as the Epidemic Type Aftershock Sequence (ETAS) model (Ogata, 1988) and the Declustering method (Reasenberg, 1985). These models utilize the detected phases and their temporal distribution to forecast the likelihood of future seismic events.

By default, GAIA uses the ETAS model as the seismicity forecasting method. The ETAS model is a statistical model that describes the temporal and spatial distribution of seismic events and is widely used in seismic monitoring. The upcoming version of GAIA will feature more advanced forecasting models, primarily based on deep learning techniques (e.g., Bi et al. (2025)). Figure 6 shows an example of GAIA's seismicity forecasting workflow, where the user uploads a known seismicity catalog and requests seismicity forecasting. GAIA Agent will then analyze the catalog, apply the selected forecasting algorithm, and return the forecasting results.

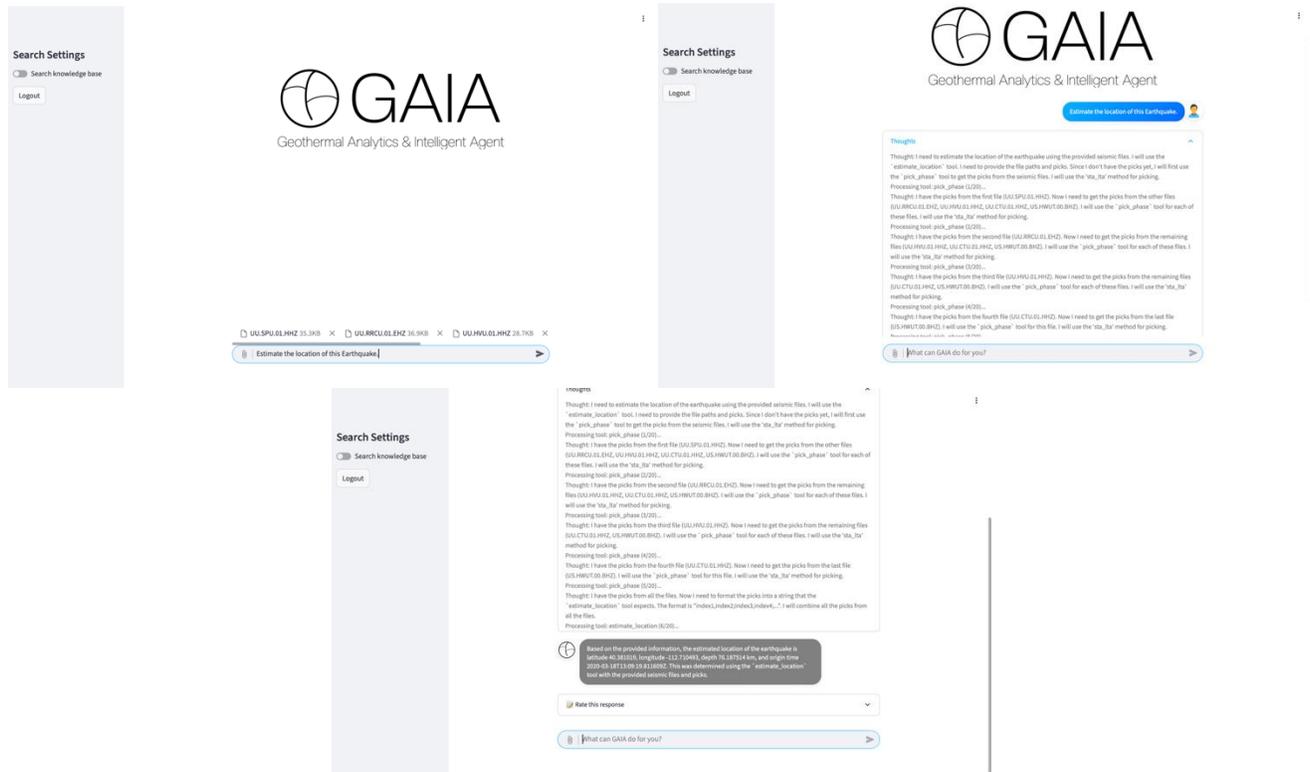


Figure 4: Example of GAIA's event location estimation workflow. The user uploads a seismic waveform file and requests event magnitude estimation (a). GAIA Agent will then analyze the waveform data and apply the selected location estimation methods (b-c).

4. CONCLUSIONS AND PLAN AHEAD

GAIA represents a pioneering application of agentic RAG workflows in geothermal field development. Unlike conventional decision-support systems, it combines autonomous task orchestration, physics-based digital twin modeling, and interactive user engagement within a single platform. This integrated approach offers a scalable pathway toward intelligent, data-driven geothermal project management and operational automation.

Due to the modular design of GAIA and the fast-evolving AI agent ecosystem, we envision several avenues for future enhancements:

- **Domain-focused LLMs:** While we currently use a general-purpose LLM (Gemma 3) as the main agent, we plan to explore domain-focused LLMs that are fine-tuned on geothermal-specific data. This could enhance GAIA's understanding of geothermal concepts and improve its ability to generate accurate and relevant responses. An example of such is GeoGPT (Huang et al., 2025a), which is a series of open weight LLMs fine-tuned on geoscience-related data.
- **MCP tools integration:** Model Context Protocol (MCP) is a standardized open-source framework that enables interaction of AI systems with external applications. By integrating MCP tools into GAIA, we can expand its capabilities to perform various geothermal data analysis tasks without needing to build the tools from scratch. Some examples of MCP servers include: QGIS MCP (<https://mcpmarket.com/server/qgis-model-context-protocol>), which is a QGIS (QGIS Development Team, 2009) MCP server, for analyzing spatial data of geothermal fields; OSDU MCP (<https://github.com/rasheedonnet/osdu-mcp>), which is Open Subsurface Data Universe (OSDU) MCP server, for accessing and managing open subsurface data stored in OSDU data platforms; and many more. Currently, there are no geothermal-specific MCP tools, but we are confident that the geothermal community will develop such tools in the near future. Therefore, we also plan to develop an MCP server for GAIA DT to allow seamless integration with other MCP-compatible applications.
- **Surrogate modeling and foundation models:** Physics-based simulations are often computationally expensive and time-consuming, which can limit their applicability in real-time geothermal monitoring and decision-making. To address this issue, we plan to incorporate surrogate modeling techniques into GAIA DT. Surrogate models are simplified representations of complex physics-based models that can approximate their behavior with significantly reduced computational costs. By integrating surrogate models into GAIA DT, we can enable faster simulations and analyses, allowing for real-time decision-making and optimization in geothermal field development. Moreover, we plan to explore the use of foundation models, which are large pre-trained models that can be fine-tuned for specific tasks. Foundation models have shown great promise in various domains, including natural language processing (NLP), computer vision (CV), and scientific computing. By leveraging foundation models in GAIA DT, we can enhance its ability to model complex geothermal systems and improve its predictive capabilities.

- Self-evolving multi-agent systems:** The current version of GAIA uses a single main agent to orchestrate the workflow and interact with users. However, as geothermal field development involves various complex tasks and processes, we plan to explore the use of self-evolving multi-agent systems in GAIA. In a multi-agent system, multiple agents can collaborate and communicate with each other to achieve common goals (Li et al., 2024). Each agent can specialize in specific tasks, such as data analysis, modeling, or visualization, and can work together to provide comprehensive solutions for geothermal field development. By implementing a self-evolving multi-agent system, we can enable GAIA to adapt and evolve its capabilities over time, allowing it to handle more complex tasks and improve its performance in geothermal project workflows.

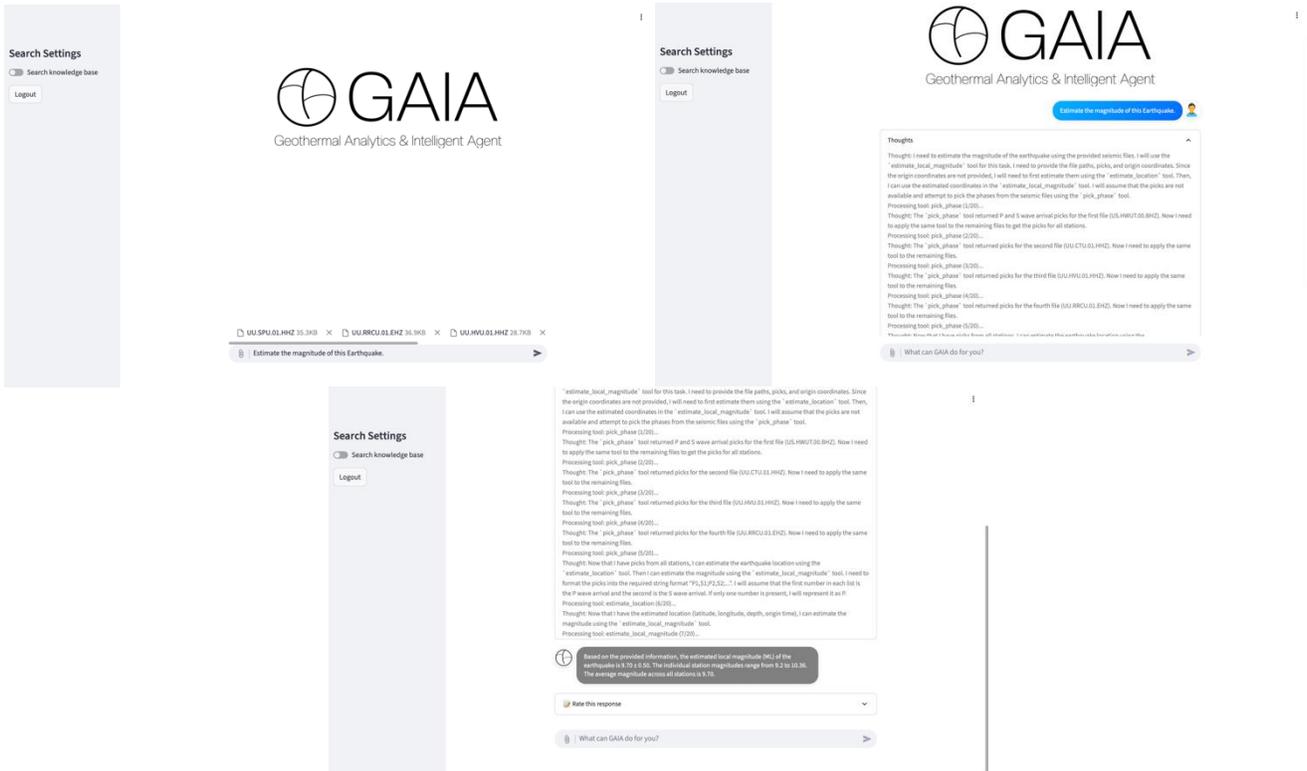


Figure 5: Example of GAIA's event magnitude estimation workflow. The user uploads a seismic waveform file and requests event magnitude estimation (a). GAIA Agent will then analyze the waveform data and apply the selected magnitude estimation methods (b-c).

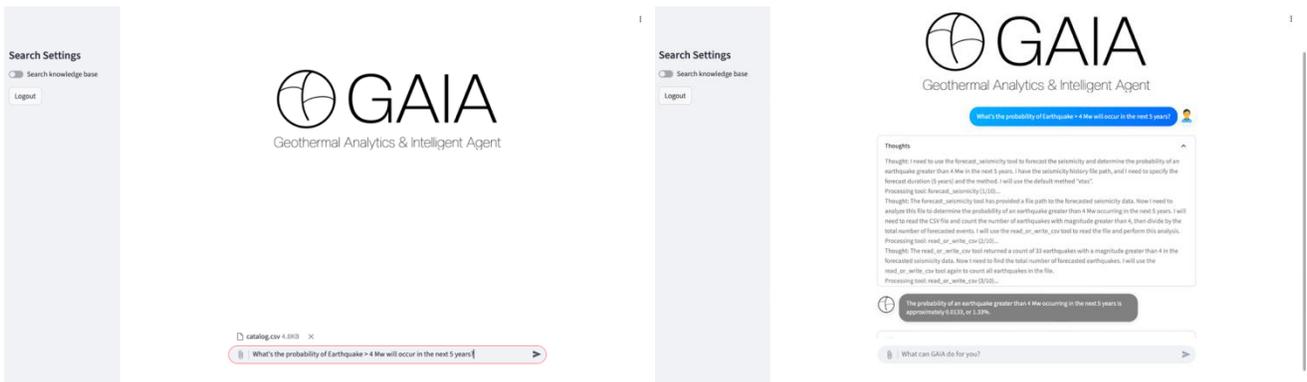


Figure 6: Example of GAIA's seismicity forecasting workflow. The user uploads a known seismicity catalog to the chatbox (a). GAIA Agent will then analyze the catalog and apply the selected forecasting model (b).

ACKNOWLEDGMENTS

The authors would like to thank the Lawrence Berkeley National Laboratory (LBNL) for providing the necessary resources and support for this research. We also acknowledge KAUST for supporting the first author's internship at LBNL, which made this work possible. We would also like to thank the Geothermal team at LBNL for their valuable feedback and contributions to the development of GAIA.

REFERENCES

- Allen, R. V. (1978). Automatic earthquake recognition and timing from single traces. *Bulletin of the seismological society of America*, 68(5):1521–1532.
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., and Wassermann, J. (2010). Obspy: A python toolbox for seismology. *Seismological Research Letters*, 81(3):530–533.
- Bi, Z., Nakata, N., Nakata, R., Ren, P., Wu, X., and Mahoney, M. W. (2025). Advancing data-driven broadband seismic wavefield simulation with multi-conditional diffusion model. *IEEE Transactions on Geoscience and Remote Sensing*, (99):1–1.
- Chan, A., Ezell, C., Kaufmann, M., Wei, K., Hammond, L., Bradley, H., Bluemke, E., Rajkumar, N., Krueger, D., Kolt, N., et al. (2024). Visibility into ai agents. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 958–973.
- Fomel, S., Sava, P., Vlad, I., Liu, Y., and Bashkardin, V. (2013). Madagascar: Open-source software project for multidimensional data analysis and reproducible computational experiments. *Journal of Open Research Software*, 1(1):e8–e8.
- Geiger, L. (1912). Probability method for the determination of earthquake epicentres from the arrival time only. *Bull. St. Louis Univ.*, 8:60.
- Gensheng, L., Xianzhi, S., Yu, S., Gaosheng, W., et al. (2024). Current status and construction scheme of smart geothermal field technology. *Petroleum Exploration and Development*, 51(4):1035–1048.
- Günther, M., Sturua, S., Akram, M. K., Mohr, I., Ungureanu, A., Wang, B., Eslami, S., Martens, S., Werk, M., Wang, N., et al. (2025). jina-embeddings-v4: Universal embeddings for multimodal multilingual retrieval. *arXiv preprint arXiv:2506.18902*.
- Gutiérrez-Oribio, D., Stathas, A., and Stefanou, I. (2025). Ai-driven approach for sustainable extraction of earth's subsurface renewable energy while minimizing seismic activity. *International Journal for Numerical and Analytical Methods in Geomechanics*, 49(4):1126–1138.
- Hanks, T. C. and Kanamori, H. (1979). A moment magnitude scale. *Journal of Geophysical Research: Solid Earth*, 84(B5):2348–2350.
- He, M., Li, Q., and Li, X. (2020). Injection-induced seismic risk management using machine learning methodology—a perspective study. *Frontiers in Earth Science*, 8:227.
- Huang, F., Wu, F., Zhang, Z., Wang, Q., Zhang, L., Boquet, G. M., and Chen, H. (2025a). Geogpt. rag technical report. *arXiv preprint arXiv:2509.09686*.
- Huang, Y., Chen, Y., Zhang, H., Li, K., Fang, M., Yang, L., Li, X., Shang, L., Xu, S., Hao, J., et al. (2025b). Deep research agents: A systematic examination and roadmap. *arXiv preprint arXiv:2506.18096*.
- Huang, Y., Gao, T., Xu, H., Zhao, Q., Song, Y., Gui, Z., Lv, T., Chen, H., Cui, L., Li, S., et al. (2025c). Peace: Empowering geologic map holistic understanding with mllms. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 3899–3908.
- Kanfar, R., Alali, A., Tonellot, T.-L., Salim, H., and Ovcharenko, O. (2025). Intelligent seismic workflows: The power of generative ai and language models. *The Leading Edge*, 44(2):142–151.
- Khorasani, M., Abdou, M., and Fernández, J. H. (2022). Web application development with streamlit. *Software Development*, 498:507.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., et al. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33:9459–9474.
- Li, X., Wang, S., Zeng, S., Wu, Y., and Yang, Y. (2024). A survey on llm-based multi-agent systems: workflow, infrastructure, and challenges. *Vicinagearth*, 1(1):9.
- Livathinos, N., Auer, C., Lysak, M., Nassar, A., Dolfi, M., Vagenas, P., Ramis, C. B., Omenetti, M., Dinkla, K., Kim, Y., et al. (2025). Docling: An efficient open-source toolkit for ai-driven document conversion. *arXiv preprint arXiv:2501.17887*.
- Lomax, A., Michelini, A., and Curtis, A. (2014). Earthquake location, direct, global-search methods. In *Encyclopedia of complexity and systems science*, pages 1–33. Springer.
- Lomax, A., Virieux, J., Volant, P., and Berge-Thierry, C. (2000). Probabilistic earthquake location in 3d and layered models: Introduction of a metropolis-gibbs method and comparison with linear locations. In *Advances in seismic event location*, pages 101–134. Springer.
- Montes, A. C., Ashok, P., and van Oort, E. Maximizing the value of geothermal data for reducing drilling risks.

- Nakata, N., Bi, Z., Qiu, H., Liu, C.-N., and Nakata, R. (2025). MI-aided induced seismicity processing and interpretation for enhanced geothermal systems. *The Leading Edge*, 44(4):265–275.
- Nakata, N., Wu, S.-M., Hopp, C., Robertson, M., and Dadi, S. (2024). Microseismicity observation and characterization at cape modern and utah forge. In Presented at the 49th Workshop on Geothermal Reservoir Engineering.
- Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association*, 83(401):9–27.
- OpenAI. Introducing deep research (<https://openai.com/index/introducing-deep-research/>).
- Peng, B., Li, C., He, P., Galley, M., and Gao, J. (2023). Instruction tuning with gpt-4. arXiv preprint arXiv:2304.03277.
- QGIS Development Team (2009). QGIS Geographic Information System. Open Source Geospatial Foundation.
- Qiu, J., Qi, X., Zhang, T., Juan, X., Guo, J., Lu, Y., Wang, Y., Yao, Z., Ren, Q., Jiang, X., et al. (2025). Alita: Generalist agent enabling scalable agentic reasoning with minimal predefinition and maximal self-evolution. arXiv preprint arXiv:2505.20286.
- Reasenber, P. (1985). Second-order moment of central california seismicity, 1969–1982. *Journal of Geophysical Research: Solid Earth*, 90(B7):5479–5495.
- Richter, C. F. (1935). An instrumental earthquake magnitude scale. *Bulletin of the seismological society of America*, 25(1):1–32.
- Shi, W., Tan, H., Kuang, C., Li, X., Ren, X., Zhang, C., Chen, H., Wang, Y., Shang, L., Yu, F., et al. (2025). Pangu deepdive: Adaptive search intensity scaling via open-web reinforcement learning. arXiv preprint arXiv:2505.24332.
- Team, G., Georgiev, P., Lei, V. I., Burnell, R., Bai, L., Gulati, A., Tanzer, G., Vincent, D., Pan, Z., Wang, S., et al. (2024). Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv preprint arXiv:2403.05530.
- Team, G., Kamath, A., Ferret, J., Pathak, S., Vieillard, N., Merhej, R., Perrin, S., Matejovicova, T., Ramé, A., Rivière, M., et al. (2025). Gemma 3 technical report. arXiv preprint arXiv:2503.19786.
- Vaghefi, S. A., Stambach, D., Muccione, V., Bingler, J., Ni, J., Kraus, M., Allen, S., Colesanti-Senni, C., Wekhof, T., Schimanski, T., et al. (2023). Chatclimate: Grounding conversational ai in climate science. *Communications Earth & Environment*, 4(1):480.
- Waldhauser, F. (2001). Hypodd-a program to compute double-difference hypocenter locations. Technical report.
- Weers, J., Podgorny, S., Taverna, N., Anderson, A., Porse, S., and Buster, G. (2024). Empowering geothermal research: The geothermal data repository’s new ai research assistant. Technical report, National Renewable Energy Laboratory (NREL), Golden, CO (United States).
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., and Cao, Y. (2023). React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*.