

Intelligenza artificiale affidabile e sovrana nelle Scienze della Terra

Dal controllo della conoscenza al controllo dell'infrastruttura

Daniele Bailo · INGV; EPOS ERIC

Geoetica e responsabilità nelle Scienze della Terra · Sala Conferenze INGV, Roma

12 ottobre 2026



ISTITUTO NAZIONALE
DI GEOFISICA E VULCANOLOGIA



Asking ChatGPT/Claude/... a geoscientific question


«According to the geological literature, what are the competing tectonic models proposed for the origin of the Gran Sasso massif, and what evidence supports each interpretation?»

According to the geological literature, what are the competing tectonic models proposed for the origin of the Gran Sasso massif, and what evidence supports each interpretation?



Elaborato per 30s

The geological literature does **not** really support a single-event origin for the Gran Sasso massif. Most authors agree that it is a **multistage structure**: a Mesozoic rifted carbonate margin was shortened and stacked during Neogene Apennine compression, and the resulting mountain belt was later dissected and further uplifted by Quaternary extension. The main controversy concerns **how much of the present structural architecture was created by thrusting versus inherited from older normal faults**, and whether some low-angle faults are contractional or extensional structures. [AGU Journals +1](#)



1. The classical out-of-sequence thrust-stack model

The classic interpretation, developed especially by G. [↓]tti & Vezzani, treats the Gran Sasso essentially as a large NE-/N-verging contractional edifice built by repeated thrusting. In this model the Gran Sasso Frontal

The answer is plausible and potentially correct.

However:

- **who selects the sources?**
- **Who receives our data?**
- **How are our prompt data used?**

+ Invia un messaggio a ChatGPT

Medio



AI in Earth Science: the problem is not only intelligence

TRUST / KNOWLEDGE


What does it know?

Where does the knowledge come from?

SOVEREIGNTY / INFRASTRUCTURE

Where do the data go?

Who controls the infrastructure?




Scientific AI is not only about model performance.

The LLM is not a scientific source

Llama2

Dataset	Sampling prop.	Epochs	Disk size
CommonCrawl	67.0%	1.10	3.3 TB
C4	15.0%	1.06	783 GB
Github	4.5%	0.64	328 GB
Wikipedia	4.5%	2.45	83 GB
Books	4.5%	2.23	85 GB
ArXiv	2.5%	1.06	92 GB
StackExchange	2.0%	1.03	78 GB

Generazione Probabilistica della Parola Successiva (Input: 'Il gatto beve il...')



Foundation model

Unclear training corpus
and information
provenance

**Probabilistic
knowledge**

Plausible, stochastic
answers, need to be
checked


Answer

Limited reproducibility,
answers follow
probabilistic patterns

Scientific AI requires governed knowledge, not only a powerful model.


The ecosystem around the LLM

Tools connect the model to the infosphere



The ecosystem around the LLM


Tools connect the model to the infosphere



We can govern the sources, tools, permissions and checks.

RockGPT: separating the model from scientific knowledge

Research papers / textbooks / trusted sources → Researcher-curated corpus



The LLM generates. Researchers govern the evidence.

Researcher-Curated Knowledge

thousands of meters of extension.
According to the geological literature, what are the competing tectonic models ...
In summary, the literature (Cardello and Doglioni [Cardello_e_Doglioni_2015.pdf](#)) does not present mutually exclusive competing models but rather a multi-stage evolutionary model: the Gran Sasso massif originated as a Mesozoic rift-related paleo-high, was subsequently folded and thrust during the Neogene (reactivating the inherited westward tilt), and is currently being dissected by Quaternary extensional faulting. The [Barchi et al. \(2021\)](#) source [Barchietal_tecphy2021.pdf](#) further supports the context of Quaternary extension in the Central Apennines, driven by sinistral transpression and normal faulting, which accommodates the stretching of the belt.

2 Sources

- 1 [Cardello_e_Doglioni_2015.pdf](#)
- 2 [Barchietal_tecphy2021.pdf](#)

24 set, 16:40

136

peer-reviewed
publications

1997–2024

SOURCE & SELECTION

Scientific journal literature,
selected by domain experts

GEOGRAPHIC FOCUS

Italy, especially
the Central Apennines


Tectonics · seismicity · seismic sequences
Geodynamics · orogenesis · geochemistry

Full corpus: Appendix B. Sources identified by author, year, journal and DOI.

RockGPT: what do we actually control?

6 models × 9 questions × 2 settings = 108 responses

Qwen2.5-72B-Instruct · Llama 3.3 70B Instruct · DeepSeek-R1 32B - Mixtral 8×7B · Gemma 2 27B · Phi-4 14B



Baseline: no corpus / RAG: curated corpus · Paired bars: retrieval / reasoning & synthesis · Model order differs

9 QUESTIONS: 6 RETRIEVAL + 3 REASONING / SYNTHESIS

2 EXPERTS · 5 CRITERIA · SCORES 0–10

Examples, shortened from Appendix A:

Which faults activated in the 1980 Irpinia sequence?

How does Moho geometry relate to Italian seismicity?

Weighted mean: content quality 30%

Coherence / relevance · logic · content

Linguistic correctness · vocabulary

Trust does not come from the model alone. It comes from architecture and governance.

From knowledge to technology governance at National and European level

Who controls the knowledge?

Sources
Selection
Validation




Who controls the technology?

Models
Infrastructure
Data access


Control over scientific sources does not automatically mean control over the AI system.

The risk of a new new lock-in

SOFTWARE ERA



AI ERA



We still have a window to avoid repeating the same dependency.


CIVITAS SINE SUFFRAGIO

30 March 2023

- ChatGPT retired after Garante della privacy measure (15M€ sanction)
- One month later, ChatGPT was reintegrated

21 March 2025

- sanction annulled in its entirety



Europe needs models and infrastructures, not only rules

«Europe has the world's most comprehensive rules for AI, yet it does not control the design and supply of the models those rules are meant to govern»

Luciano Floridi, 2026

RULES

Set of rules for European Jurisdiction and beyond

GDPR
AI Act
Data Act
Digital regulation

MODEL SOVEREIGNTY

Develop models we can adapt and maintain

Build on genuinely open-source models
Use distillation and further training
Retain the ability to improve and release models

DEPLOYMENT SOVEREIGNTY

Run systems under European control

Use infrastructure under European jurisdiction and operational control
Govern access to prompts, documents and logs
Retain control over service operation and continuity

What could this mean for a scientific institution?

01

LOCAL / INSTITUTIONAL AI

Sensitive documents · institutional knowledge

RAG · internal assistants · experimentation · Local LLM

02

SHARED PUBLIC INFRASTRUCTURE

Research infrastructures · national HPC / cloud

European compute · shared scientific AI services

03

COMMERCIAL FRONTIER SERVICES

Tasks requiring frontier capabilities · low-risk workloads

External services when advantageous

**A sovereign institution does not need one model.
It needs an architecture that keeps alternatives open.**

From Researcher-Curated Knowledge to Sovereign Scientific AI

Research Institutions

TRUSTWORTHY AI KNOWLEDGE

corpus
provenance
verification
Local LLM
LLM fine tuning

National HPC providers

SOVEREIGN AI AI ECOSYSTEM

Research Institutions
And HPC provider
Collaborating for
→
A SOVEREIGN AI

LLM models training
Inference Services
High level skills

Trustworthy scientific AI requires governance of knowledge.
Sovereign scientific AI extends that governance to the infrastructure itself.

Europe is not starting from scratch

Examples discussed by Luciano Floridi (2026)

MODELS

EuroLLM

Apertus

ALIA

Amália

Mistral*

Development capability

Pan-European

Switzerland

Spain

Portugal

France

PROGRAMMES

OpenEuroLLM

SOOFI

FAIR

Shared research efforts

European consortium

Germany

Italy

INFRASTRUCTURE

EuroHPC

CINECA / Leonardo

AI Factories

Compute and expertise

European HPC ecosystem

Italian computing capacity

AI compute and services

We should aim at shared inference services for public institutions

Reliable access, data protection, support and skills

1. Scientific knowledge is a strategic asset.

2. The infrastructure used to interpret it should be treated as one too.

3. Europe and Italy are still on time to tackle the challenge

4. Geoethics should include technological capabilities issues.



The European Open-LLM Landscape

NON-EXHAUSTIVE · WORK IN PROGRESS

Who trains their own model, and who releases everything.
A snapshot of a moving target — groups are missing.

FULLY OPEN

weights + data + code

OWN MODEL
pretrained from scratch

OWN MODEL

6

-  **BSC** Salamandra
-  **Swiss AI** Apertus
-  **PleIAs**
-  **OpenLLM-France** Luciole
-  **Sapienza NLP** Minerva
-  **OpenEuroLLM**

OPEN WEIGHTS

weights public, data/code not

OWN MODEL

5

-  **EuroLLM**
-  **OpenGPT-X** Teuken
-  **Danish Foundation Models** Mimir
-  **Mistral AI** Mistral
-  **AI Sweden** GPT-SW3

POST-TRAINING ONLY
on a foreign base model

POST-TRAINING ONLY

4

-  **Occiglot** multi-country
-  **CroissantLLM**
-  **Soofi** Soofi S
-  **Lapa LLM**

POST-TRAINING ONLY

16

-  **Language Technology Group** NorMistral
-  **LightOn AI** PAGnol
-  **SpeakLeash** Bielik
-  **NB AI Lab** Borealis
-  **BritLLM**
-  **TurkuNLP** gpt3-finnish
-  **ILSP** Krikri
-  **jpacifico** Chocolatine
-  **LINAGORA Labs** Labess
-  **Slayer Lab** Slayer
-  **PLLuM**
-  **INSAIT** BgGPT
-  **CJVT** GaMS
-  **Aleph-Alpha**
-  **Unbabel** Tower-Plus
-  **BottleCap AI** ThinkingCap

<https://lnkd.in/p/dBSvv7qY>

References

RESEARCH PAPERS

[Bailo et al. \(2026\)](#)

Researcher-Curated knowledge governance for trustworthy scientific AI in the geosciences: the rockGPT approach

Earth Science Informatics 19, 152.

[Floridi \(2026\)](#)

Rules Without Models: Why Europe Must Build the AI It Regulates

Philosophy & Technology 39, 162.

[Lewis et al. \(2020\)](#)

Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks

NeurIPS 33. arXiv:2005.11401.

[Hinton, Vinyals & Dean \(2015\)](#)

Distilling the Knowledge in a Neural Network

METHODS & GUIDANCE

[Anthropic \(2024\)](#)

Building effective agents

Engineering article, 19 December 2024.

[Anthropic \(2025\)](#)

How we built our multi-agent research system

Engineering article, 13 June 2025.

[UNESCO \(2021\)](#)

Recommendation on the Ethics of Artificial Intelligence

Adopted in 2021; publication edition 2022.

[Open Source Initiative](#)

The Open Source AI Definition – 1.0

EUROPEAN INITIATIVES

[Albert API \(DINUM\)](#)

[L'Assistant \(DINUM\)](#)

[F13: project website](#)

[Svea: impact report 2025](#)

[EuropAI: cooperation project](#)

[GPT-NL: project website](#)

[Delos: product website](#)

