GENESIS Open Call

Technical Guidelines

Funded by
the European Union

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission. Neither the European Union nor the granting authority can be held responsible for them. The GRAIL project has received funding from the Horizon Europe programme under grant agreement No. 101298421

Table of contents

1. Introduction 3
2. Data Strategy and Sovereignty 3
3. GRAIL Industrial Use Cases 4
4. Core Architecture Tiers 4
5. Key Platform Capabilities Relevant to the Applicants 6
    5.1 Expected Role of External Solutions 6
    5.2 Integration Touchpoints 7
    5.3 Complementarity and Non-Overlap Requirement 7

1. Introduction

Welcome to the GRAIL Open Call Technical Guidelines. For Genesis (OC1), this document provides high-level context on GRAIL and the technical expectations for data collection, acquisition tools and benchmarking. It should be read together with the Genesis Open Call Terms and Conditions and their eight-topic annex; those call documents define the binding application requirements.

Genesis applicants should focus on the requirements and expected outputs of their chosen topic. The GRAIL Foundational Robotics Model (GFRM), the five industrial Use Cases, and the Digital Platform are being developed progressively; using them is not a general prerequisite for participation in OC1. Specific guidance for Ascent and Summit will be provided separately. For any questions, please contact us through the official GRAIL Open Call communication channels indicated in the Terms and Conditions document.

Good luck!

The GRAIL Consortium

2. Data Strategy and Sovereignty

GRAIL aims to develop a unified multimodal data backbone combining industrial partner data provided under data-sharing agreements, repurposed open corpora (e.g. Open X-Embodiment) and prior EU projects, and new data acquired through Simulation, Use Cases, Open Calls, and Testing Hubs. Genesis (OC1) seeks complementary data assets, acquisition methods and evaluation resources as defined in the eight topics. Applicants should describe their proposed data modalities, volumes, quality controls, provenance, annotation and access arrangements; the detailed obligations are those set out in the Terms and Conditions and Sub-Grant Agreement.

To balance openness with industrial confidentiality, GRAIL applies a tiered private / protected / public data governance model: public assets (benchmarks, lightweight models) support open dissemination and reproducibility, while sensitive industrial datasets remain under controlled, sovereignty-preserving access. GRAIL is developing a dedicated Data Broker to support controlled, sovereignty-preserving access to datasets. Access to specific services, datasets or interfaces will depend on their availability and the applicable agreements. Genesis applicants should not assume access to these resources when preparing their proposals.

GRAIL will validate its research through industrial Use Cases and Testing Hubs. These are useful project context, but the Genesis topics are not restricted to the GRAIL industrial Use Cases, and applicants are not expected to have access to those sites unless expressly agreed under their selected topic.

3. GRAIL Industrial Use Cases

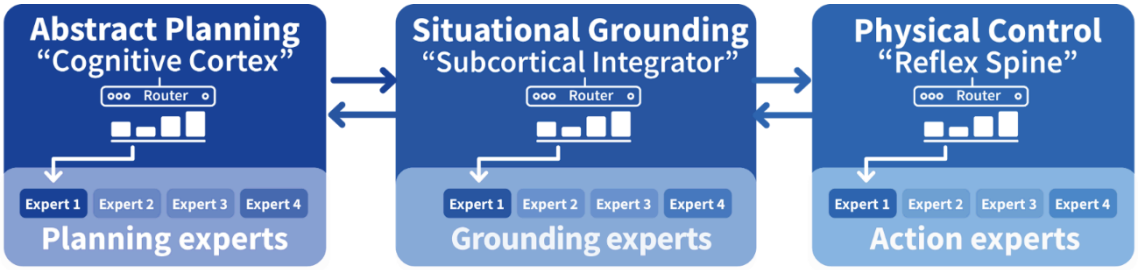
GRAIL includes five industrial Use Cases and four GRAIL Testing Hubs. These resources will become available progressively according to the project roadmap. For Genesis, the eight published topics define the relevant tasks and deliverables; access to a specific industrial Use Case or Testing Hub is not a general eligibility requirement.

At least 40 FSTP (Financial Support to Third Parties) projects are planned across the three Open Calls, supporting data collection, fine-tuning, and industrial adoption of the GFRM at different maturity levels.

Genesis contributions focus primarily on data, acquisition tools, demonstrators and evaluation methods. These may support the later development of the GFRM, but OC1 beneficiaries are not generally required to integrate with a GFRM release or the Digital Platform.

4. Core Architecture Tiers

The GRAIL architecture is conceived as an agentic, hierarchical, and modular Generative-AI-for-Robotics ecosystem, designed to unify foundation-model intelligence with real-time, safety-critical robotic execution. It responds to the structural challenge known as the robotics “data gap”: unlike text and vision, robot experience cannot be scaled through brute-force data collection, as it is slow, costly, and infeasible at trillion-scale. Bridging this gap requires modularity, systematic knowledge reuse, and targeted data generation, rather than monolithic, end-to-end approaches that are data-inefficient, difficult to explain, and unable to cope with industrial complexity.



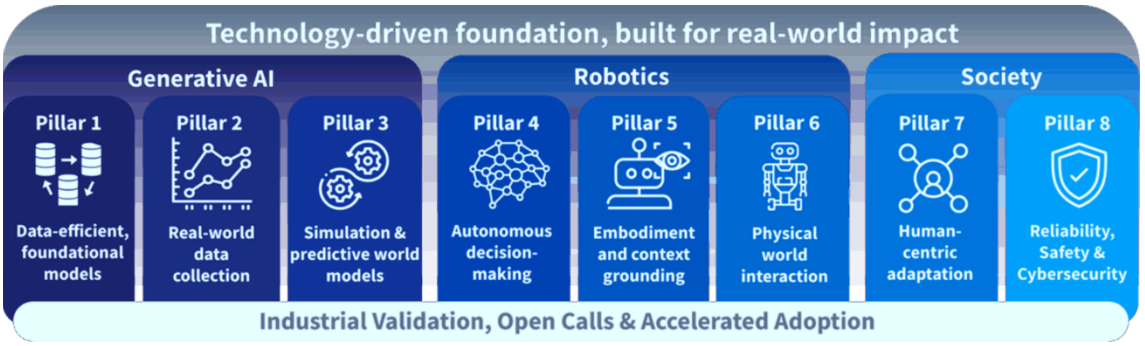
At the heart of GRAIL lies the GRAIL Foundational Robotics Model (GFRM): a novel architecture that unifies Generative AI and Robotics into a coherent framework for

general-purpose, safe, and explainable robot intelligence. The GFRM is not a monolithic black box but a hierarchical, distributed, and agentic system built on a Mixture-of-Experts (MoE) paradigm. Inspired by dual-system cognitive theories — distinguishing fast, intuitive processing from slower, deliberative reasoning — GRAIL separates high-level symbolic reasoning, mid-level perceptual grounding, and low-level motor control into three specialised agents, each operating at a distinct temporal and semantic scale:

- **Abstract Planning Agent ("Cognitive Cortex") —** translates human intent and complex industrial tasks into structured, verifiable, and auditable symbolic plans.
- **Grounding Agent ("Subcortical Integrator") —** adapts abstract plans to the physical world, embedding them into embodiment constraints, perceptual context and real-time affordances.
- **Physical Control Agent ("Reflex Spine") —** executes grounded plans in real time, reacting to disturbances through reflexive policies, motor primitives, or reinforcement-learned skills.

The GFRM is designed to exchange information between high-level planning, perceptual grounding and low-level control, with feedback improving decisions across layers. This is an evolving project architecture, not an interface requirement for Genesis beneficiaries.

GRAIL is also developing shared state representations, agent communication and predictive modelling to support future integration. OC1 proposals should not assume these capabilities will be available when their projects start.



The eight GRAIL scientific pillars encompass foundation-model development, data, simulation, robotics and responsible deployment. The associated research is carried out by the consortium across the project lifetime; Genesis focuses on the contributions described in its eight topics:

1. **Generative AI —** foundational models, data engines, simulation, and lifelong learning (Pillars 1–3).
2. **Robotics —** control, perception, planning, and real-world interaction (Pillars 4–6).

3. **Societal integration —** human-centricity, ethics, trust, and safety (Pillars 7–8).

These complementary pillars will progressively supply the technical foundations for later GRAIL releases and industrial validation.

5. Key Platform Capabilities Relevant to the Applicants

GRAIL is progressively developing an agentic Generative-AI-for-Robotics platform, including foundational model components, data infrastructure, interfaces and validation facilities. These components should not be assumed to be operational or accessible when Genesis projects begin.

Technical expectations differ across the three Open Calls. For Genesis, applicants should propose feasible ways to collect, document, assess and, where contractually required, share outputs relevant to their topic. Integration with GRAIL APIs, models or Testing Hubs is not a universal prerequisite; such dependencies must be explicitly agreed and matched to actual availability.

Depending on the selected topic and the resources available during implementation, possible interactions with GRAIL include:

- Accessing curated datasets via the tiered private/protected/public data broker. Where such datasets are actually available and relevant under the applicable access rules.
- Where appropriate and available, using documented GRAIL interfaces or reference assets, subject to the corresponding access and licensing arrangements.
- Where explicitly agreed in the Individual Mentoring Plan, validating data-collection tools or benchmarks in a Testing Hub or another suitable environment; access to GRAIL industrial pilots is not generally required for Genesis.
- Making agreed datasets, annotations, tools or benchmark outputs available to GRAIL under the applicable contractual access rights, data protection and licensing provisions.

GRAIL's platform and interfaces are intended to support interoperability and later reuse of successful outputs. For Genesis, applicants should prioritise portable, well-documented deliverables without assuming access to unreleased models, services or internal infrastructure.

5.1 Expected Role of External Solutions

Genesis applicants should address one of the eight published OC1 topics. Depending on the topic, relevant outputs may include:

- Dexterous manipulation, long-horizon robot execution and tool-use demonstrations.
- Egocentric video and multimodal/contact-rich capture relevant to robotic tasks.
- Structured planning, failure-diagnosis and recovery traces.
- Documented and reproducible data-capture and annotation methods.
- Benchmarking kits, stress-test conditions and evaluation protocols.

The technical suitability of a Genesis proposal is assessed against its chosen topic and the Open Call documentation, not against a general requirement to consume or modify GRAIL's foundational model.

5.2 Integration Touchpoints

Technical coordination with GRAIL may involve the following, where relevant to the selected Genesis topic and agreed in the project plan:

- The tiered data broker (private / protected / public) for dataset access and contribution, rather than direct modification of GRAIL's internal data pipelines.
- The GRAIL Testing Hubs, using portable integration kits and standardised benchmark scripts for reproducible evaluation.
- The Individual Mentoring Plan and Milestone Review process, which formally tracks technical progress and integration status.

The intended reuse of Genesis outputs, including potential access by GRAIL partners and later Open Call beneficiaries, must follow the specific rights and obligations established in the Sub-Grant Agreement. Applicants should plan for lawful, documented data access and licensing, taking into account confidentiality, personal data and third-party rights. This does not imply that future beneficiaries are obliged to use any particular dataset or that access is granted outside the agreed contractual terms.

5.3 Complementarity and Non-Overlap Requirement

GRAIL will deliver core platform capabilities, including:

- The GRAIL Foundational Robotics Model (Planning, Grounding, and Control agents).

- The Integrated Situational Model and latent communication substrate.
- The tiered private/protected/public data engine and data broker.
- The GRAIL Digital Platform (API Gateway, configuration registry, model/artefact repository, CI/RobOps pipelines).
- Reliability, safety, and cybersecurity frameworks across all three agents.
- The four GRAIL Testing Hubs and the harmonised evaluation/KPI framework.

Three-Phase Development Roadmap

GRAIL follows a progressive development roadmap across there phases (*Genesis*, *Ascent* and *Summit*), advancing from early research and data collection to industrial validation and deployment. The technical capabilities available to beneficiaries will depend on the maturity of the project at each stage. Specific requirements and access arrangements will be defined through the corresponding Open Call documentation and applicable agreements.

Genesis builds early GFRM capabilities: it trains the first Alpha release on limited, controlled, general datasets, bootstraps generative simulation assets and synthetic data, validates on simulated and lab robots, and sets up data and compute governance. Ascent strengthens reliability by combining generative AI with rule-based and constraint-driven safeguards, adds continual learning, knowledge transfer and stronger Sim2Real, and tests the Beta release in controlled industrial pilots. Summit fine-tunes the model into a stable MVP release with real-world adaptability, explainability and human-in-the-loop learning, and scales deployments with major manufacturing partners.

Technical Baseline Available at Each Open Call

Different Open Calls take place at different stages of GRAIL development. The technical requirements for Genesis are established by its own topic descriptions and contractual documents.For Genesis, proposals should respond to their chosen topic and avoid assuming that the emerging GRAIL platform must be used or modified. In particular, applicants:

- may propose independent but well-documented acquisition, annotation or benchmarking methods relevant to their selected topic.
- should describe data quality, provenance, interoperability, reproducibility and the intended rights of access and reuse.
- must identify any reliance on GRAIL models, APIs, data resources or Testing Hubs and assess availability with the mentoring team.
- should follow the applicable technical, ethical, safety and data-protection requirements in the Open Call documents.

Examples of potentially relevant Genesis contributions include:

- Novel data-capture methods and multimodal demonstration datasets for robotic tasks.
- Robot and egocentric data collection, structured annotations and practical capture tools.
- Reproducible benchmarking suites, stress-test conditions and robust evaluation protocols.
- Well-documented data assets and tools that can support subsequent GRAIL research subject to agreed licensing and access conditions.