LORENTZ CENTER WORKSHOP

Accelerating Astrophysical Discovery

with Foundation Models

Participant booklet

14–18 September 2026 · Leiden, The Netherlands

VERSION 1.0 · 11 AUGUST 2026

WELCOME

A week built for making

Dear Lorentz Workshop participant,

In a little over a month, we will meet in Leiden for an interactive week devoted to pushing the bounds of how we study the cosmos. We want the workshop to turn diverse expertise into shared decisions, concrete artifacts, and collaborations that continue after Friday.

This booklet introduces the consortium's ambition, explains the initial working set of four Key Problem Areas (KPAs), lays out the working timetable, and suggests a few things to think about before you arrive. You do not need polished proposals. Bring concrete examples, informed questions, and a willingness to build across disciplinary boundaries.

AT A GLANCE

Workshop	Accelerating Astrophysical Discovery with Foundation Models
Dates	Monday 14 to Friday 18 September 2026
Venue	Lorentz Center@omega, Leiden, The Netherlands
Start times	Monday at 10:00; Tuesday–Friday at 9:30

The working style is closer to an open-source project hackathon than a sequence of lectures: short framing talks, focused working sessions, plenary synthesis, and a named deliverable each day.

WORKSHOP COMMUNITY

List of participants

Last name	First name	Affiliation
Albert	Joshua G.	California Institute of Technology
Albin	Thomas	Freelancer Scientist
Angeloudi	Eirini	University of Hamburg, Germany
Bester	Hertzog Landman	SARAO, South Africa
Bhihe	Cedric	BSC, Spain
Bulmus	Turan	Google
Callingham	Joseph	ASTRON/UvA
Caron	Sascha	Radboud University
Cheng	Ting-Yun	University of Groningen
Ciuca	Ioana	Stanford University, USA
Cruz	Marcos	Instituto de Física de Cantabria, Spain
Cuesta Lazaro	Carolina	Flatiron Institute / New York University
De la Torre Luque	Pedro	Universidad Autónoma de Madrid (UAM)
Eckner	Christopher	Instituto de Astrofísica de Canarias
Fernández Sopuerta	Carlos	ICE/CSIC, Spain
Glombitza	Jonas	FAU, Germany
Heng	Ik Siong	University of Glasgow, UK
Huertas-Company	Marc	IAC
Huppenkothen	Daniela	University of Amsterdam, The Netherlands
Jain	Divesh	University of Heidelberg
Keenan	Ryan	Max Planck Institute for Astronomy
Keller	Christoph	National Solar Observatory
Khan	Sebastian	Starling Bank
Leonhard	David	University of Bonn
Lopez	Melissa	Nikhef
Loukili	Soulaymane	BSC, Spain
Madireddy	Sandeep	Simons SkAI, Argonne National Laboratory
Pérez Romero	Judit	IFIC (UV/CSIC)
Ruiz de Austri	Roberto	Universidad de Valencia
Samia	Noelle	Simons SkAI
Smirnov	Oleg	Rhodes University & SARAO
Sotoudeh	Hadi	University of Cambridge
Stevance	Héloïse	Oxford University, UK
Tallada Crespí	Pau	CIEMAT-PIC
Ting	Yuan-Sen	Ohio State University, USA
Vacheret	Antonin	Laboratoire Corpusculaire de Caen UMR 6534, Caen, France
Villaescusa-Navarro	Francisco	Simons Foundation
Weigel	Philip	Harvard University
Zaharijas	Gabrijela	University of Nova Gorica
Zhao	Yue	SURF, NL

Consortium fields

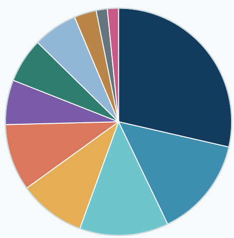

	Share of field entries
AI	**28.6%**
Cosmology	**14.3%**
Radio	**12.7%**
Computational physics	**9.5%**
Gravitational waves	**9.5%**
Gamma-rays	**6.3%**
Neutrinos	**6.3%**
Optical	**6.3%**
X-Ray	**3.2%**
CRs	**1.6%**
Infrared	**1.6%**

THE CONSORTIUM

Two connected pillars

PILLAR 01

Design science for human–AI collaboration

The first goal is to design how future astrophysicists and cosmologists will do science in collaboration with AI. This is more than adding chat interfaces to existing tools. It means redesigning workflows so humans, reasoning language models, simulators, archives, instruments, and foundation models can jointly form hypotheses, retrieve evidence, run checks, update beliefs, and communicate results.

PILLAR 02

Build a joint representation of physical data

The second goal is to build the technical substrate: a model that learns across observations, simulations, synthetic observations, language, code, and instrumental context. It should support meaningful embeddings for retrieval and serendipity, along with generative capabilities for conditional inference across linked views of the same physical system.

One shared ambition. The two pillars meet in an agentic discovery system: reasoning models guide scientific work, the joint-representation model grounds that work in physical evidence, and humans remain responsible for judgement.

KEY PROBLEM AREAS

KPAs 1 and 2

The four KPAs are a starting framework, not a fixed list. A central workshop goal is to test whether they frame the right scientific problems—and to refine, combine, replace, or add KPAs as our discussions and use cases demand.

1 Serendipity through embeddings

A physically meaningful embedding space can surface rare objects, surprising analogues, counterexamples, and neighbouring evidence a scientist may not know to request.

Workshop question Can it find scientifically meaningful candidates faster and more reliably than current workflows?

Evidence of progress Positive controls, rare-class retrieval, synthetic anomalies, baseline comparisons, and expert review.

2 Generative inference across instruments

Instruments see the universe through different resolutions, cadences, noise, coverage, and selection effects. Conditional generation may connect those views through super-resolution, deconfusion, gap filling, missing-modality prediction, or cross-instrument reconstruction.

Workshop question Is the generated view supported by evidence, rather than merely plausible-looking?

Evidence of progress Controlled degradation, held-out modalities, calibrated uncertainty, forward-model checks, and classical baselines.

Scientific standard: useful generative inference must expose what is observed, what is assumed, what remains uncertain, and how the output can be checked.

KEY PROBLEM AREAS

KPAs 3 and 4

3 Observations to physical states

Multimodal observations could constrain candidate physical states or simulation initial conditions. Candidate states are evolved through system and instrument simulators, then compared back to observations.
Workshop question Do forward runs remain physically admissible and observationally consistent?
Evidence of progress Forward-model consistency, stable physical structure, residual analysis, uncertainty, and an auditable chain of assumptions.

4 Human–machine discovery

This is the workflow lens: how should scientists, reasoning language models, tools, simulators, archives, and the joint-representation model form, test, revise, and communicate hypotheses together?
Workshop question Does the combined workflow improve search, triage, hypothesis generation, validation, or communication without hiding uncertainty or responsibility?

Nodes are physical data; relevancy links become the training-time attention mask.

THE RELEVANCY GRAPH

We envision a large graph database whose nodes are physical data and whose links encode relevancy. Unlike text, whose sequence yields a causal mask, physical data have no universal order. During self-supervised training, sampled graph connectivity becomes the attention mask. It defines a more general attentional field—cross-modal, many-to-many, and hierarchical—within which the model predicts, reconstructs, or aligns withheld data.

BEFORE LEIDEN

Bring ingredients, not finished answers

ASTROPHYSICS + COMPUTATIONAL PHYSICS

Think about use cases where AI could help humans study the cosmos and do better science if it could understand, reason about, and generate physical data. Prepare a list of astrophysical data you can access. Any modality is welcome, including:

- observations, catalogues, spectra, images, time series, and event lists;
- simulation code, snapshots, initial conditions, and synthetic observations;
- multimodal observations of the same systems;
- instrument responses, calibration context, uncertainties, and provenance.

AI + COMPUTATIONAL INFRASTRUCTURE

Think about the architectural challenges facing a joint-representation model that must learn deep, self-supervised representations of large physical datasets. Also consider the infrastructure required to train across physical datasets that can dwarf the data used to train language models:

- long context, heterogeneous tokenisation, and shared latent spaces;
- compute-to-data training across distributed archives and HPC systems;
- data access, preprocessing, provenance, monitoring, and validation;
- staged scaling, simulator-backed synthetic data, and failure modes.

If you span both groups, bring both perspectives. If you span neither exactly, bring the constraints, examples, and questions from your own domain.

THE WEEK

The week at a glance

Day	Focus	Main deliverable
Monday	Vocabulary, use cases, evaluation	Shared Vocabulary and Discovery-Evaluation Map v0
Tuesday	Data, representation, provenance	Relevancy Graph v0 for selected use cases
Wednesday	Architecture and training	Architecture and Training Protocol v0
Thursday	MVP and hackathon	MVP Implementation Plan and Hackathon Artifacts v0
Friday	Roadmap, governance, funding	Consortium Roadmap and Funding Package v0

THE HAND-OFF

Monday chooses what matters. Tuesday defines the relationships the data and system must respect. Wednesday turns those constraints into a model and training path. Thursday turns that design into workstreams and artifacts. Friday gives the work owners, resources, governance, and continuity.

Working rhythm Each day combines brief framing, active work, plenary synthesis, and a concrete closeout. Tuesday–Thursday begin by reviewing overnight AI work; Monday–Wednesday end by agreeing what the AI should prepare for the next morning. Coffee breaks and lunches are part of the programme: use them to connect ideas across disciplines.

Formal working hours are 10:00–16:30 on Monday, 9:30–17:00 Tuesday–Thursday, and 9:30–13:30 on Friday. The building is open for arrivals and coffee from 09:00 on Monday.

MONDAY · 10:00–16:30

Use cases and discovery workflows

Deliverable Shared Vocabulary and Discovery-Evaluation Map v0

Time	Session
09:00–10:00	**Arrival and coffee** The formal programme begins at 10:00.
10:00–10:30	**Kickoff talk** AI-first science, the two pillars, and KPAs as a starting framework.
10:30–12:00	**Science 2.0 + use cases** Work from long-term visions back to concrete workflows and candidate use cases.
12:00–13:30	**Lunch**
13:30–14:30	**Plenary synthesis** Share, refine, and select a small set of flagship use cases.
14:30–15:00	**Coffee break**
15:00–16:15	**Define evaluation targets** Controls, baselines, anomaly tests, forward checks, uncertainty, provenance, and human review.
16:15–16:30	**Closeout + AI brief** Confirm cases, evaluations, owners, and the overnight synthesis for Tuesday.
16:30 onward	**Welcome reception** Informal continuation after the formal programme.

Keep asking What would count as evidence that this workflow produces better science, rather than simply more output?

TUESDAY · 9:30–17:00

Data, representation, and provenance

Deliverable Relevancy Graph v0 for Monday's selected use cases

Time	Session
9:30–10:00	**Kickoff + overnight review** Correct Monday's synthesis; introduce the graph and tokenisation challenges.
10:00–10:30	**Map data objects** Use participants' lists to connect cases to observations, simulations, instruments, code, and provenance.
10:30–11:00	**Coffee break**
11:00–12:00	**Graph working session** Define nodes, relevancy links, metadata, provenance, ambiguities, and training relationships.
12:00–13:30	**Lunch**
13:30–14:30	**Plenary graph synthesis** Share group results and assemble Relevancy Graph v0.
14:30–15:00	**Coffee break**
15:00–16:30	**Tokenisation + constraints** Record what representations must preserve and the constraints for Wednesday.
16:30–17:00	**Closeout + AI brief** Confirm owners; commission a draft schema or prototype and architecture primer.

Keep asking What relationship makes these objects scientifically relevant to one another, and what provenance must survive every transformation?

WEDNESDAY · 9:30–17:00

Long-context architecture and training

Deliverable Architecture and Training Protocol v0

Time	Session
9:30–10:00	**Kickoff + overnight review** Review the primer and graph work; frame long-context, recurrent-depth, and graph-attention challenges.
10:00–10:30	**Translate graph constraints** Preserve typed objects, relationships, units, uncertainty, and metadata.
10:30–11:00	**Coffee break**
11:00–12:00	**Architecture working session** Use AI to sketch candidate architecture families, objectives, and training hypotheses.
12:00–13:30	**Lunch**
13:30–14:30	**Plenary synthesis** Compare designs, identify common components, and retain justified alternatives.
14:30–15:00	**Coffee break**
15:00–16:30	**Training + prototype brief** Set objectives, simulator feedback, deployment, validation, scaling, risks, and Thursday artifacts.
16:30–17:00	**Closeout + AI brief** Confirm hypotheses, protocol, owners, and overnight prototype work.
17:00–20:00	**Workshop dinner** Informal continuation after the formal programme.

THURSDAY · 9:30–17:00

MVP architecture and hackathon

Deliverable MVP Implementation Plan and Hackathon Artifacts v0

Time	Session
9:30–10:00	**Kickoff + overnight review** Review the prototype, state the MVP target, and share AI-assisted coding approaches.
10:00–10:30	**Self-organise workstreams** Choose concrete artifacts and owners across technical, documentation, and roadmap work.
10:30–11:00	**Coffee break**
11:00–12:00	**Open hackathon I** Build toy-scale artifacts in the main room or breakouts.
12:00–13:30	**Lunch**
13:30–14:30	**Open hackathon II** Continue building; demonstrate useful progress as it appears.
14:30–15:00	**Coffee break**
15:00–16:15	**Show-and-tell + integration** Demonstrate artifacts, connect them to use cases and evaluation, and identify gaps.
16:15–17:00	**Closeout + Friday handoff** Confirm artifacts, issues, owners, next work, and Friday’s resource questions.

Thursday rule: every artifact should connect to a use case, a model capability, and an evaluation target.

FRIDAY · 9:30–13:30

Roadmap, governance, and funding

Deliverable Consortium Roadmap and Funding Package v0

Time	Session
9:30–10:00	**Funding + continuity kickoff** Review the week's artifacts and frame resources, partnerships, deployment, and maintenance.
10:00–10:30	**Map resource needs** Compute, data, simulators, archives, expertise, funding, and missing roles.
10:30–11:00	**Coffee break**
11:00–11:30	**Funding and partnerships** Identify opportunities, partners, institutional support, and lead owners.
11:30–12:00	**Commitments + roadmap** Record desired participation; confirm milestones, owners, cadence, and invitees.
12:00–13:30	**Closing lunch** Final confirmations and departures.

PRACTICAL MATTERS

Discounted hotel reservations are limited, so please book as soon as possible if you need accommodation.
For practical questions or help with accommodation, contact Maria Krebbers at acceleratingastrophysical@lorentzcenter.nl.

accelerating-astrophysical-discovery.github.io

Bring a use case. Bring a constraint. Leave with something we can build together.