

KYMA — Learning abstractions for checkable AI reasoning

Public research overview · Draft v1 · 2 October 2026

Miroslav Šotek · Marbach SG, Switzerland · [ORCID 0009-0009-3560-0851](#) · protoscience@anulum.li

Application stage: preparing an application to EIC Pathfinder Challenges 2026, DeepRAP. The proposal has not been submitted; no grant has been awarded. The proposed research duration is 36 months.

The research question

Can a learned abstraction transfer to unfamiliar combinations of tasks while remaining independently checkable and affordable to compute?

KYMA investigates phase representations in coupled oscillator systems. Relative phases provide a candidate way to represent relations, while learned synchronisation motifs are tested as reusable abstractions. Specialised language-model interfaces would translate eligible language and numerical inputs into typed predicates, tool requests and plans. A separate checker would assess defined claims before an output is accepted.

The checker would certify only the stated properties under its assumptions. It would not guarantee that every statement is true or that observational evidence establishes causation. Transfer, reliability and any benefit from oscillator dynamics remain questions to test.

How the proposal would test the idea

Conventional and phase-based representations would use matched sources, tools, checking and resource budgets. Separate analytic controls would help distinguish useful phase structure from the contribution of integrated dynamics. Held-out compositions and changes in task conditions would test whether the learned representations transfer beyond the training setting.

The proposed laboratory pilots concern scientific structure discovery from dynamical-system trajectories and bounded decision support over dated regulatory documents. Evaluation would report incorrect admissions, refusals, failed tasks and negative findings as well as successful outputs. Cost, latency and measured local energy would be assessed per correctly completed task, including retries and checking; unmetered hosted-service energy would remain unknown.

A bounded SC-NeuroCore FPGA investigation is proposed for selected kernels. It is an implementation comparison, with its continuation dependent on evidence. It is not a claim that an FPGA will train or replace a frontier language model.

Planned outputs and collaboration

Planned outputs include experimental protocols, controlled comparisons and ablations, typed interfaces, independent checkers, reproducible benchmarks, publications and rights-compatible model or adapter releases. Existing software in our research ecosystem provides starting components; integration and validation would be part of the proposed work.

We welcome discussion with researchers interested in compositional abstraction, formal checking, scientific machine learning, bounded planning and measurement of computational resources. Scientific collaboration, doctoral supervision and recruitment would require separate agreements; no institution or funded position is currently committed through this overview.

[DeepRAP official call and guidance](#) · [Contact Anulum](#)

Full scientific presentation: anulum.li/research/programmes/kyma/