NEURA PARSE

RESEARCH AND INNOVATION OVERVIEW / FOR RESEARCH PARTNERS

Research that ends in a result someone else can check.

Neura Parse runs one applied robotics programme (NODERIQ), publishes quantum research (QANTIS, qmesh, QMANN), and keeps its engineering record public.

ROBOTICS PROGRAMME

QUANTUM RESEARCH

OPEN SOURCE

WHAT IS WHAT

We keep four kinds of work clearly separate.

Saying exactly what stage each is at is part of the method.

01

PRODUCTS

Commercial products

NowFlow, NeuraOS, and QFlow Studio are products we sell, each with a different job.

02

PROGRAMME

Applied research programme

NODERIQ is a research programme on robots working together, not a product yet.

03

RESEARCH + OSS

Public research and open source

QANTIS, qmesh, QMANN, our datasets, and side projects are all public, with their tests and limits.

HOW WE DO RESEARCH

A claim moves forward only when the evidence does.

Every project follows the same four steps, and each step has to be passed in order.

01

Frame

State the decision, the setting, the uncertainty, and the limits.

02

Baseline

Set the classical or current method as the bar before testing a new one.

03

Exercise

Run repeatable cases and record versions, routes, costs, and failures.

04

Decide

Stop, refine, keep watching, or promote, against agreed gates.

APPLIED RESEARCH PROGRAMME / NODERIQ

Can many machines still work together when they see the world differently?

NODERIQ studies robot work orders that can be checked, a shared picture of the world that admits doubt, mixed fleets, human authority, and working without a cloud link.

RESEARCH PROGRAMME WORKING TOWARDS A PRODUCT. NOT A RELEASED PRODUCT OR FIELD DEPLOYMENT.

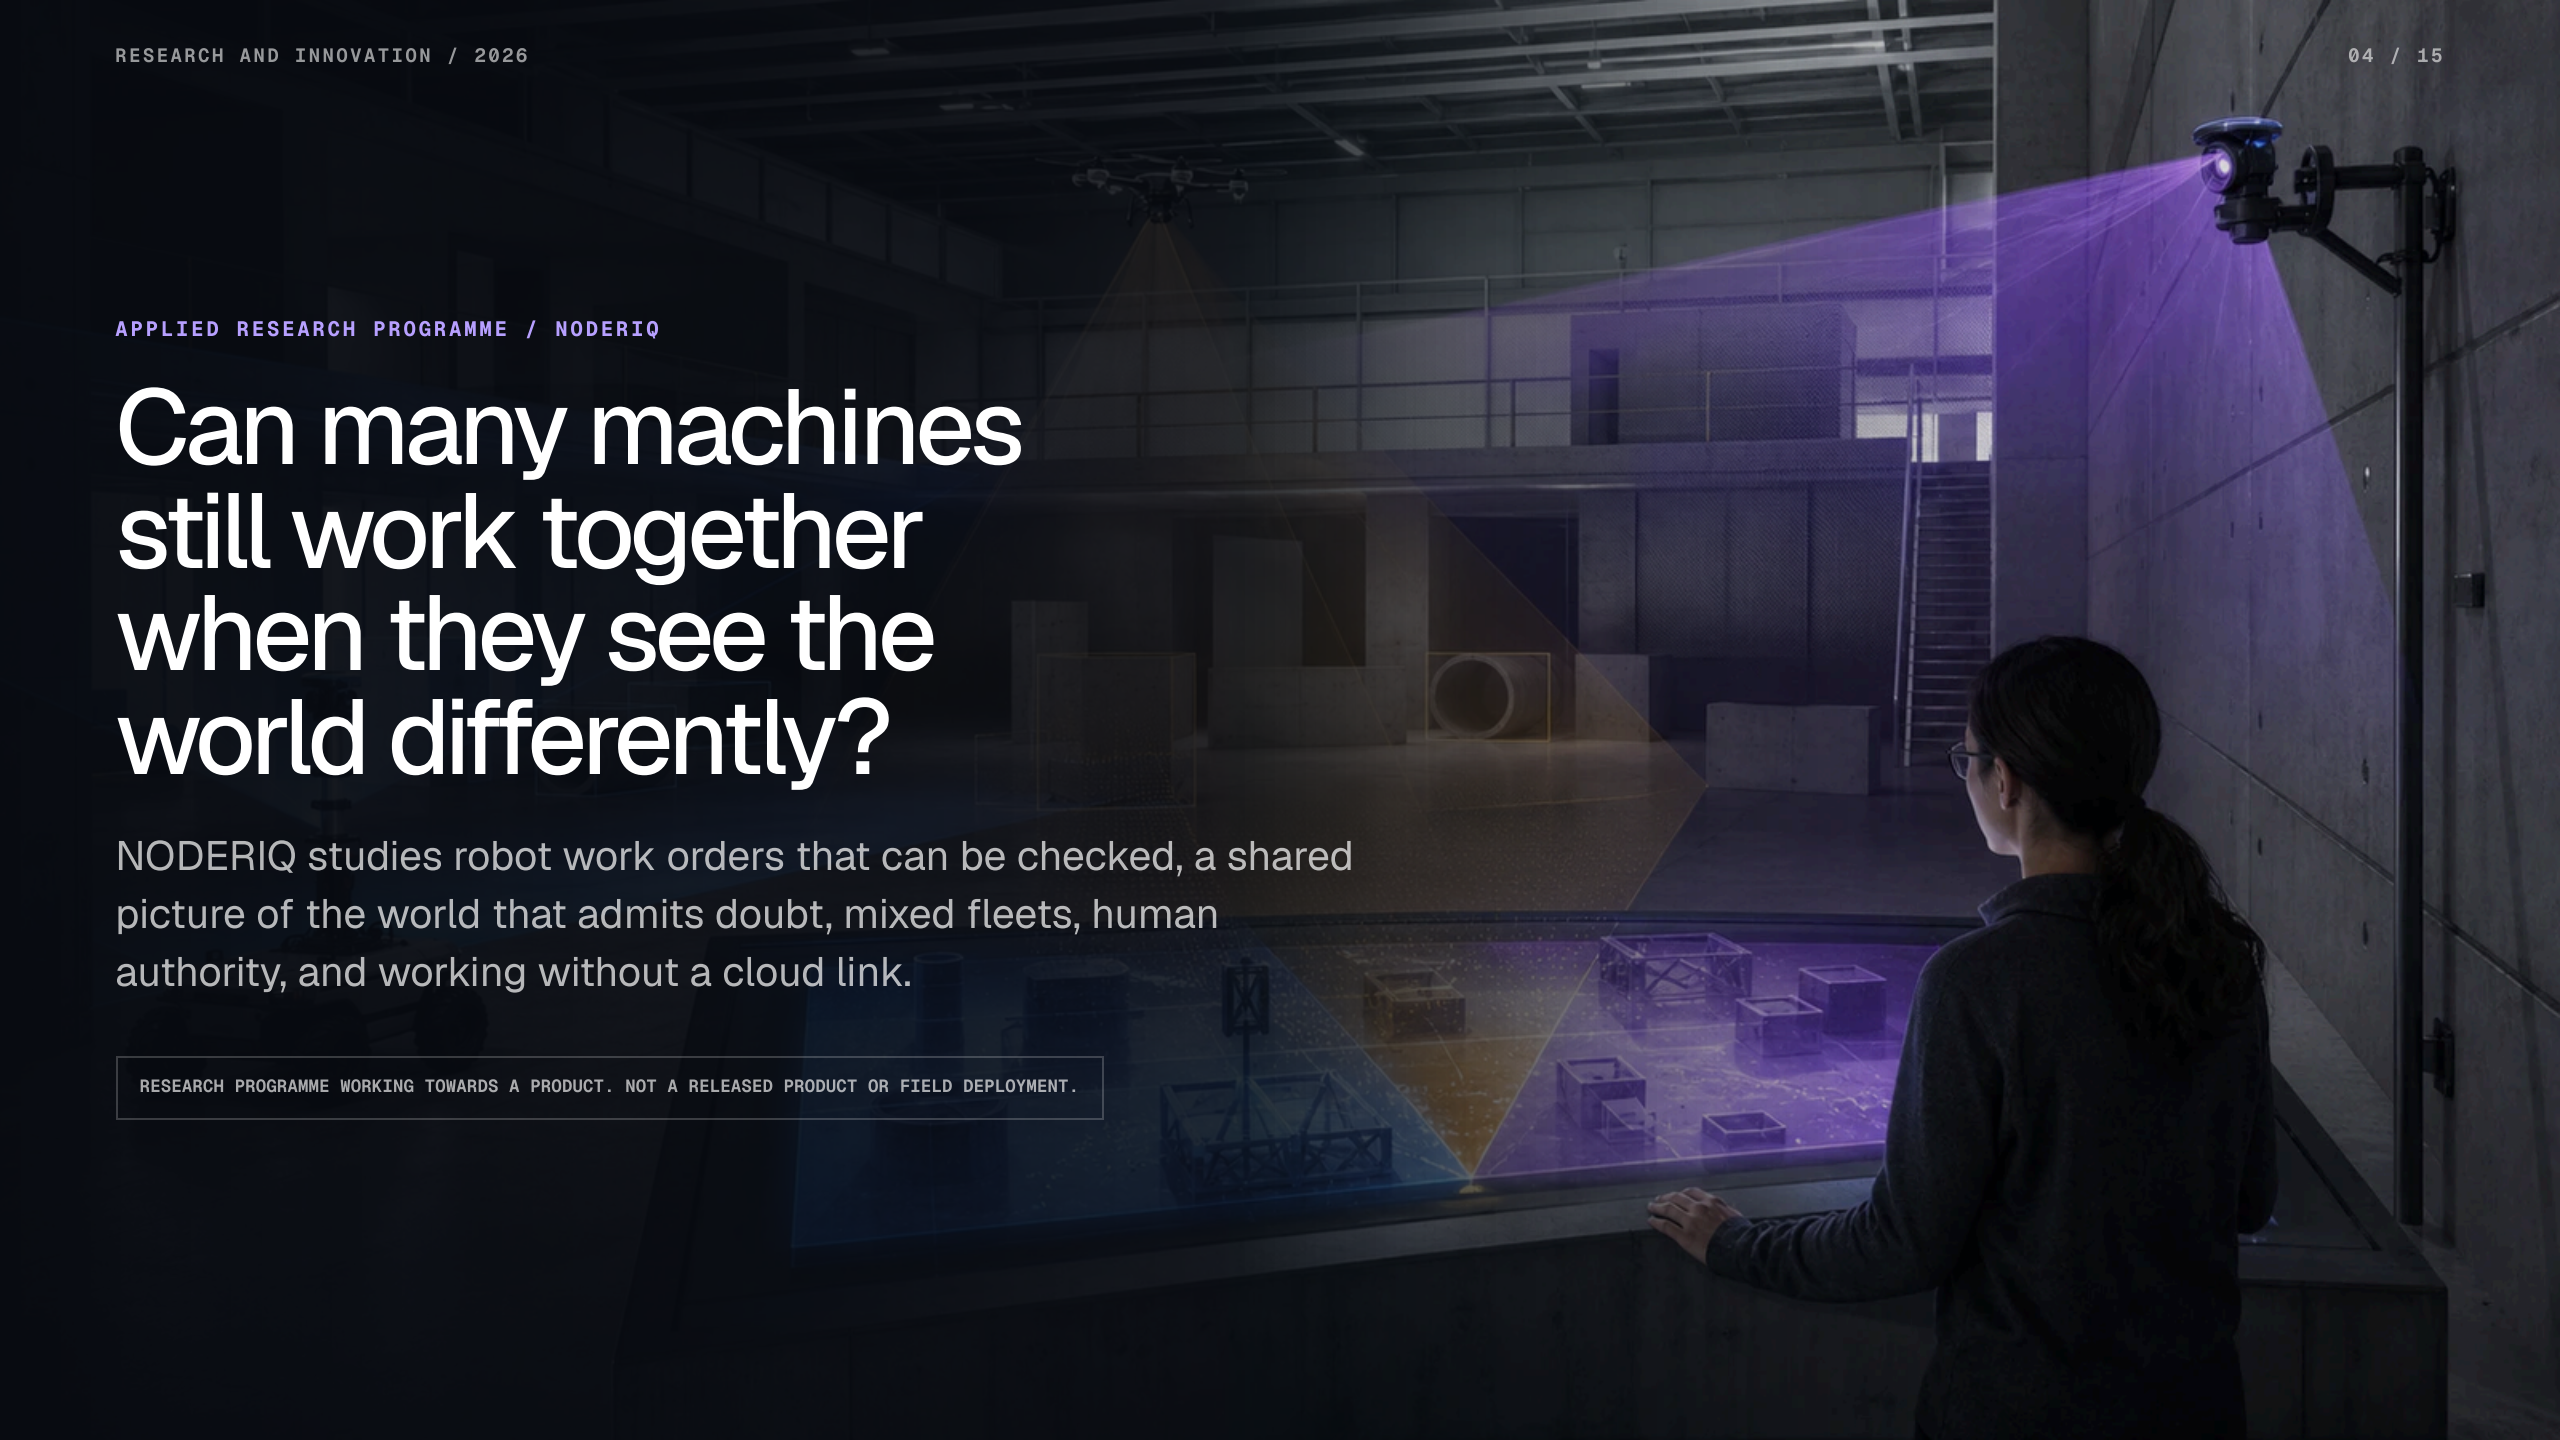

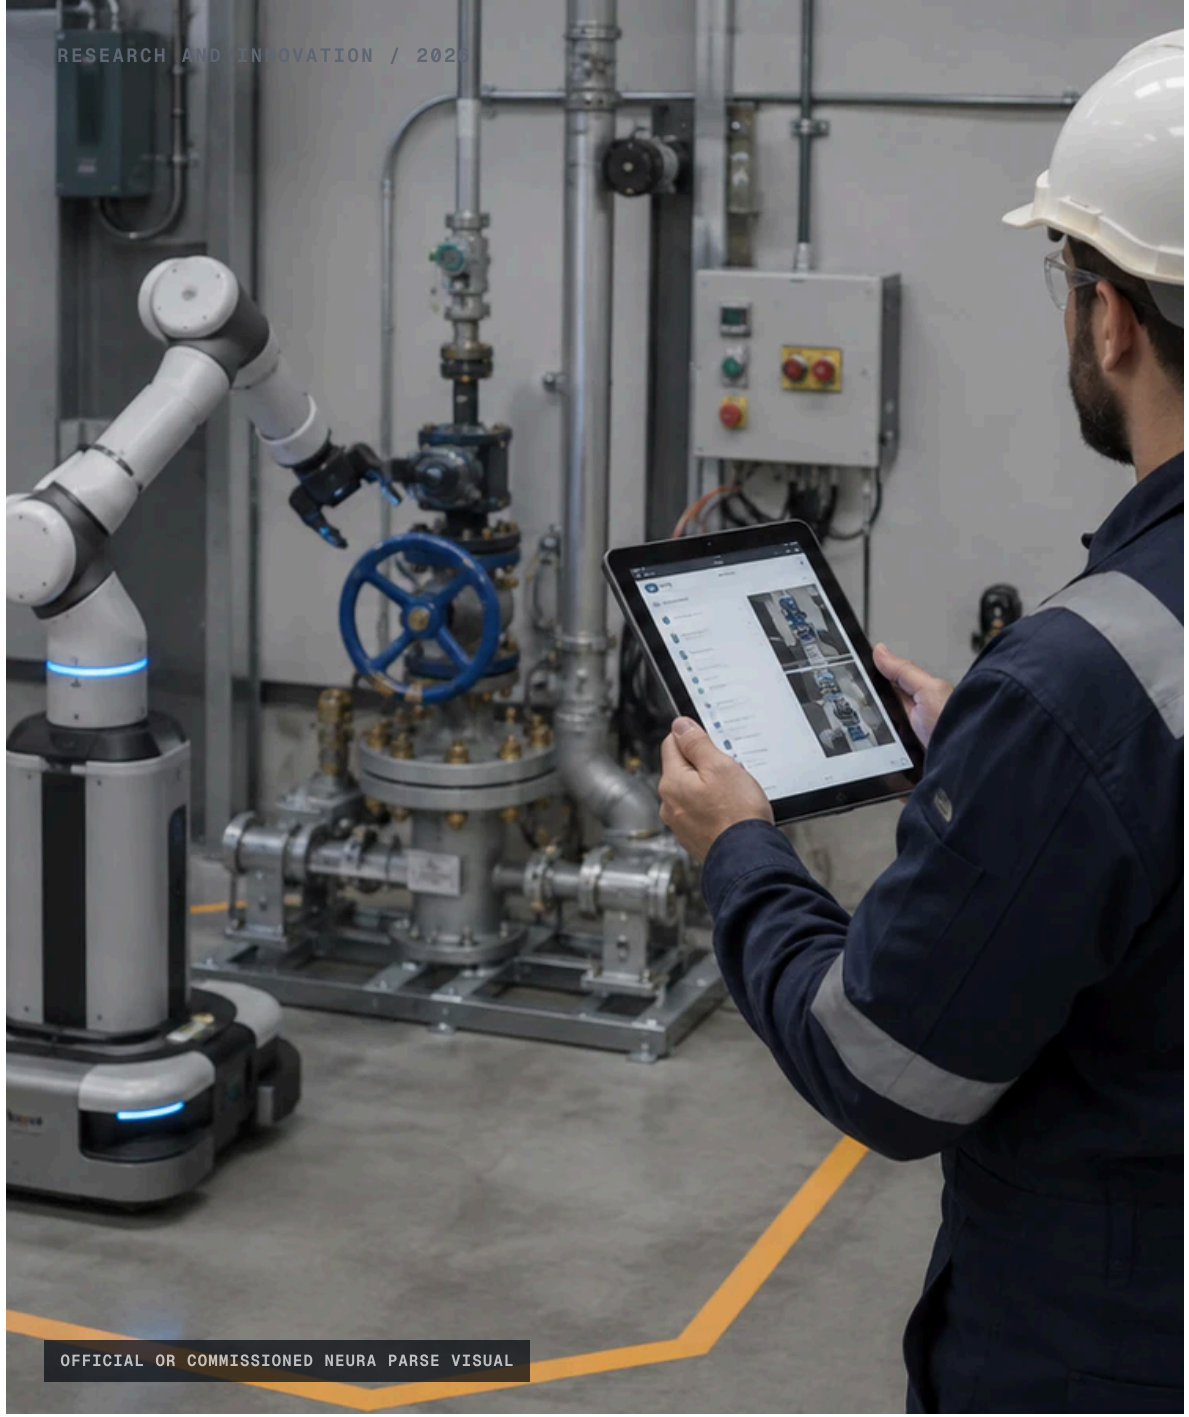

OFFICIAL OR COMMISSIONED NEURA PARSE VISUAL

NODERIQ / A ROBOT WORK ORDER

A spoken instruction is only the start of a robot work order.

The proposed path matches the task to a capable machine, keeps the machine's own safety, records progress and completion, and can replay it all.

CONTRACT	**Define the job** The goal, the test for done, the resources, the owner, and the safe stop.
EXECUTION	**The machine keeps its say** The fleet manager, the robot, and its controller keep their own safety rules.
CLOSEOUT	**Prove what happened** Progress, exceptions, human interventions, recovery, and proof of completion can be replayed.

OFFICIAL OR COMMISSIONED NEURA PARSE VISUAL

NODERIQ / WHERE QUANTUM FITS

Quantum methods have to beat the classical baseline to be used.

We try quantum or quantum-inspired methods only on planning or estimation tasks that can wait and are never safety-critical.

BASELINE	**Classical by default**	The core must work with ordinary computers and without a constant cloud link.
GATE	**Measure the trade-off**	Compare quality, uncertainty, run time, cost, repeatability, and fit for the job.
PROMOTION	**No advantage claims**	A good research result never widens what a safety-critical system may do.

QANTIS

QUANTUM DECISION RESEARCH / QANTIS

Can a quantum computer help a system decide under uncertainty?

QANTIS is our research on quantum-assisted decisions: planning when you are unsure of the state, and matching sensor tracks to targets.

01 INFER	02 RISK	03 OPTIMISE	04 VERIFY

Planning and acting stay on classical computers. The papers make no claim about total run time.

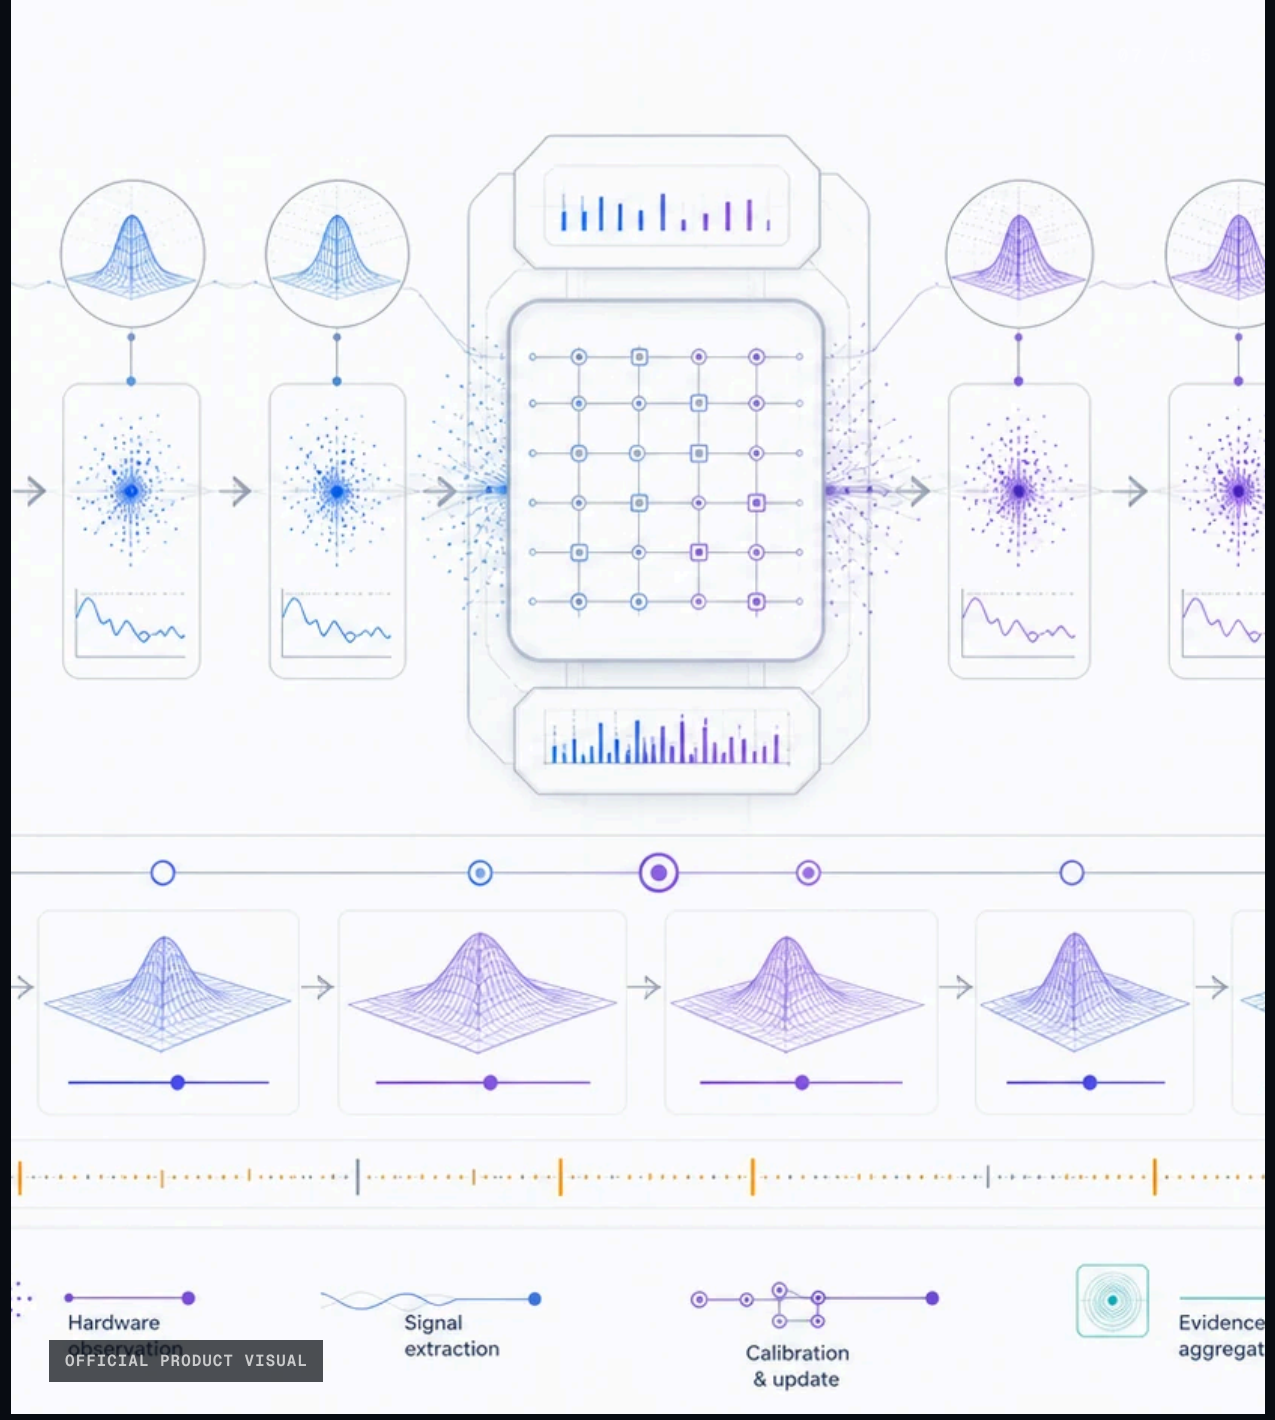

QANTIS / THE TWO PAPERS

Two public papers, and what they do and do not show.

The latest paper uses an IBM Heron processor for one step: updating a belief. Neura Parse's IBM Partner Plus membership is company-level only, not an endorsement.

2603.00785	**First paper, Feb 2026**	Planning under uncertainty and multi-target tracking
2607.06760	**Second paper, Jul 2026**	8-step and 12-step main runs, 20-step and 32-step controls
DECISION CHECK	**Matches exact Bayes**	The chosen action agreed with exact Bayes in every reported check
BOUNDARY	**Planning stays classical**	No comparison of total run time is made
ACCESS	**Two editions**	MIT Community Edition is public, Collaborator Edition is private

OFFICIAL OR COMMISSIONED NEURA PARSE VISUAL

OPEN-SOURCE INFRASTRUCTURE / QMESH

One shared format for quantum programs, with signed records.

qmesh is an open-source layer that sits above Qiskit, Cirq, or PennyLane. It gives programs one typed format, schedules hybrid jobs, and signs the record.

IR	**One format, many machines** Gate, neutral-atom, photonic, and pulse programs can share the same typed format.
PROVENANCE	**Signed records** Records are signed and hash-chained, so they can be verified offline.
STATUS	**Apache-2.0, 157 tests** Works with Qiskit, Cirq, and PennyLane. It does not replace them.

OPEN-SOURCE LIBRARY / QMANN

One learning library, three ways to run it.

QMANN is an open-source neural network library with a quantum memory layer. It estimates the cost before any paid hardware run.

01

THEORETICAL

Free / unlimited qubits

Study the architecture on paper, with no qubit limit.

02

SIMULATION

Free / up to 20 qubits

Run the model with noise on a simulator, repeatably.

03

HARDWARE

Paid / 4-12 qubits

See the cost estimate first, then run on IBM Quantum or IonQ.

OFFICIAL OR COMMISSIONED NEURA PARSE VISUAL

WHERE OUR PRODUCT HELPS / QFLOW STUDIO

Research travels better when the whole workflow is visible.

QFlow Studio, our web product, keeps the brief, the circuit, the generated code, the run, the results, and the reviewer's verdict in one place.

AUTHOR	**Design** The visual circuit and its Qiskit, Cirq, or OpenQASM code stay linked.
RUN	**Run** Local simulators and real hardware routes are labelled by their status.
HANDOFF	**Hand over** Assumptions, outputs, failed runs, limits, and the next decision are kept.

WHAT IS PUBLIC

Everything public, so a collaborator has somewhere to start.

Datasets, code, tests, papers, and engineering notes are public, each with its status marked.

17 **Public datasets** Hugging Face, CC-BY-4.0, as of 4 Jul 2026	**157** **qmesh tests** Passing, open-source repository
2 **QANTIS papers** Public on arXiv, 2026	**3** **Open-source side projects** TaskNebula, NeuraBar, OrbIDE preview

HOW TO WORK WITH US

Pick the kind of collaboration that fits how mature the question is.

A paper, a pilot, and a consortium each end with a different kind of evidence.

ROUTE	GOOD FOR	WHAT YOU HAVE AT THE END
Research collaboration	A method, a benchmark, a paper	Question, baseline, protocol, result, limits, next hypothesis
NODERIQ programme work	Mixed robot fleets, poor connectivity, human authority	Interface map, test cases, exception replay, go or no-go
Small product pilot	One real workflow or experiment	An acceptance record and a stop, adjust, or expand decision
Consortium formation	Several partners designing work packages	Roles, IP, data, integration, assurance, exploitation, governance

CONSORTIUM DESIGN / ANY CALL

A good consortium gives each partner a checkable job.

The same structure works for any funding call: name the real problem, assign the interfaces, test against baselines, and plan how results get used.

01	USE CASE	**Problem owner**	Describes the real setting, the decision, the constraints, and what success looks like.
02	TECHNOLOGY	**System and research partners**	Bring the models, workflows, devices, networks, quantum methods, and adapters.
03	ASSURANCE	**Independent reviewer**	Sets the tests, the safety and security limits, human factors, and review criteria.
04	EXPLOITATION	**Route to use**	Owns IP, data rights, the path to a product, standards, papers, and the route to market.

NEURA PARSE

NEXT STEP

Bring a hard question and its classical baseline.

We will pick the right project to work in, agree what evidence the result needs, and keep research, programme, and product claims separate.

NEXT

neuraparse.com/research

qflow.studio / github.com/neuraparse