Promoting
Knowledge Management Culture

AI in Agricultural Research: What's in it for NARO?

Dr. Olimi will draw on his experience to show what AI can achieve when integrated into a research workflow. A discussion session will focus on what research opportunities can be piloted, how to test AI tools responsibly, what safeguards are needed for data protection, skills required, and when human judgement must remain in control.



Expedito Olimi, PhD
Erwin Schrödinger Research Fellow
University of Southampton

22 July, 2026 **12pm**

zoom

ID: 872 8536 2662

Passcode: 396443

Register Now

info@naro.go.ug

www.naro.go.ug

Artificial Intelligence: Domesticating the technology capabilities for NARO

Expedito Olimi (PhD) · Erwin Schrödinger Research Fellow

What to expect?

THE TALK

- What is AI?
- Where it can fit in the research lifecycle
- A case application in microbiome research
- **AI sovereignty: Institutional or enterprise-level choice and owning models**

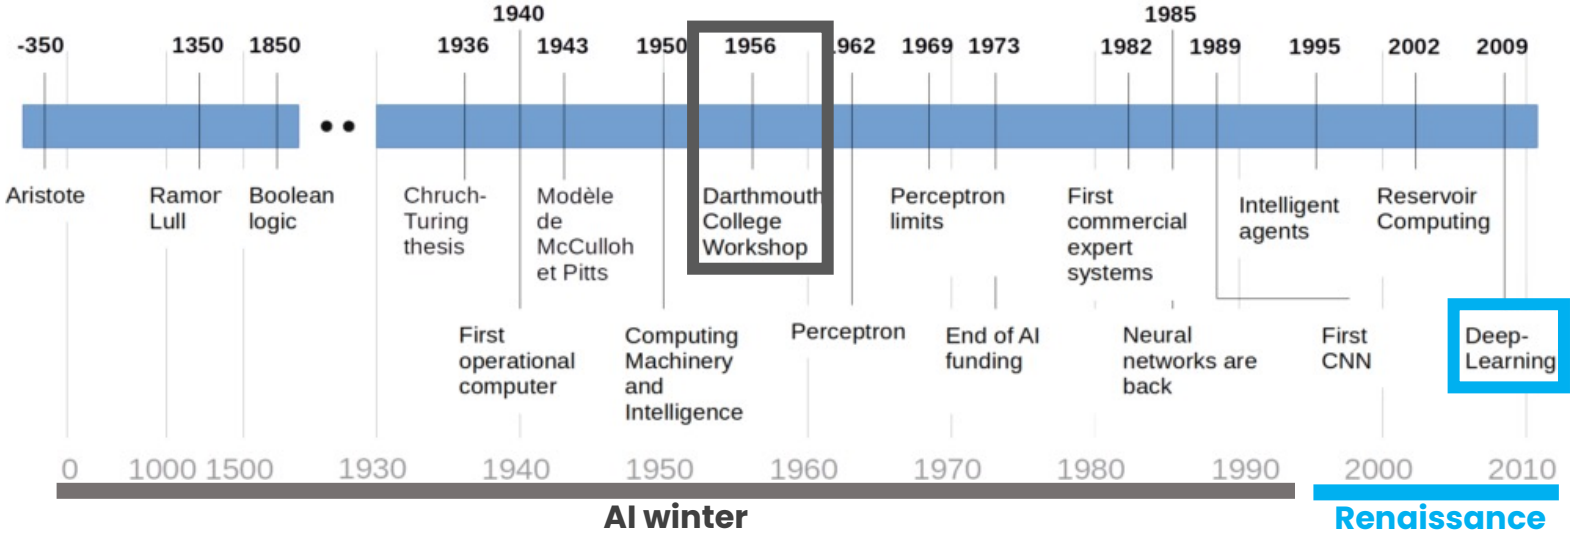
THE DISCUSSION

- Which research opportunities to pilot
- How to test AI tools responsibly
- Safeguards for data protection
- The skills teams will need
- When human judgement must stay in control

Goal: a primer towards strategic integration of AI use in institutional workflow **(safely & responsibly)**


2015
SELF-DRIVING CARS IN 2 YEARS
2016
RADIOLOGISTS OBSOLOLETE IN 5 YEARS
2026
RADIOLOGISTS IN REGULAR CARS DRIVING TO WORK

AI is not new





1950s

John McCarthy

Coined the term “artificial intelligence.” at Dartmouth college during a summer school



1959

Arthur Samuel

Built the first program that learned by itself, playing checkers. Coined the term **Machine learning**



Today

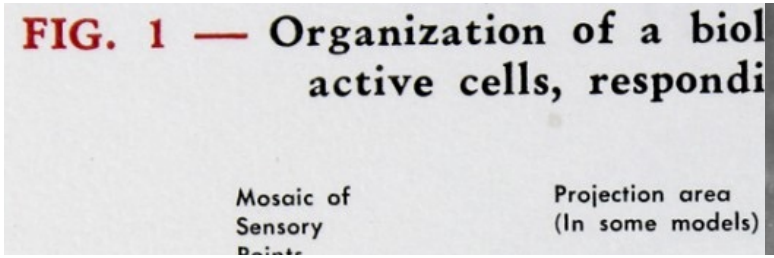
All around us

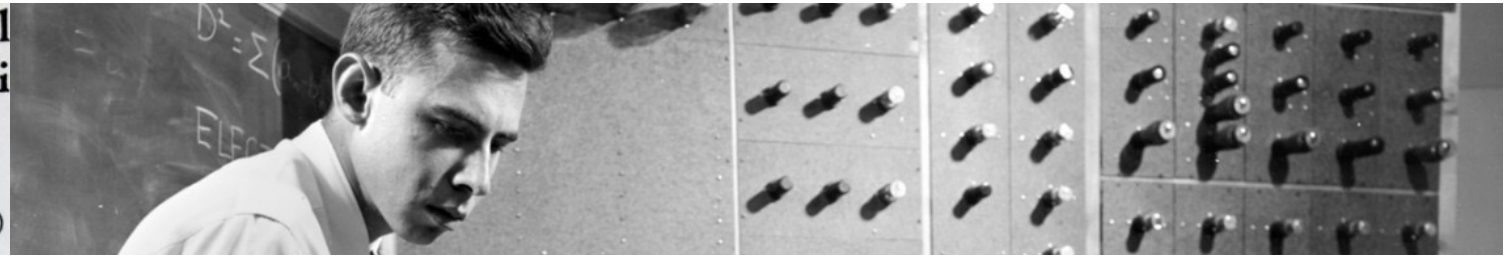
Fraud checks, online shopping and self-driving-car sensors.

The technology is over 70 years old . We are just in the renaissance age of AI.

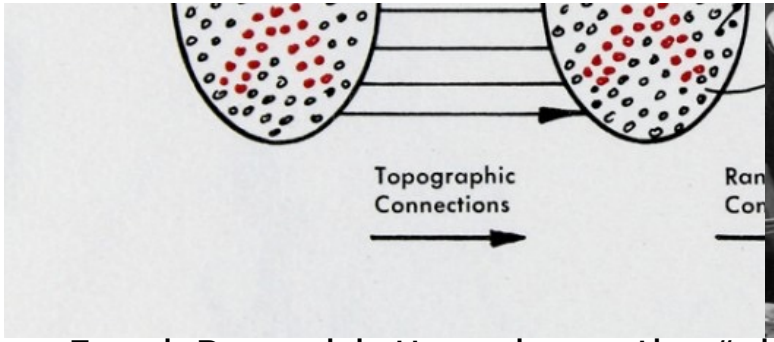
Bipartisan Policy Center, “AI: Facts and Myths” (2026); He, Cao & Tan (2025), NSR 12(5): nwaf050.

The Mark 1 perceptron: a mechanical prototype of a biological neurone in action





New York Times 1958: "the Navy [has] revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence."

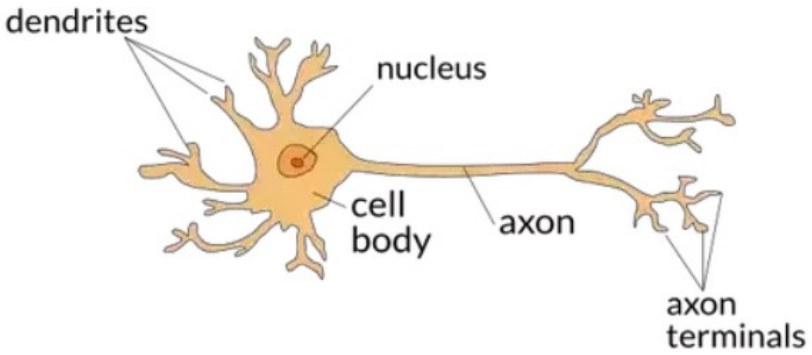


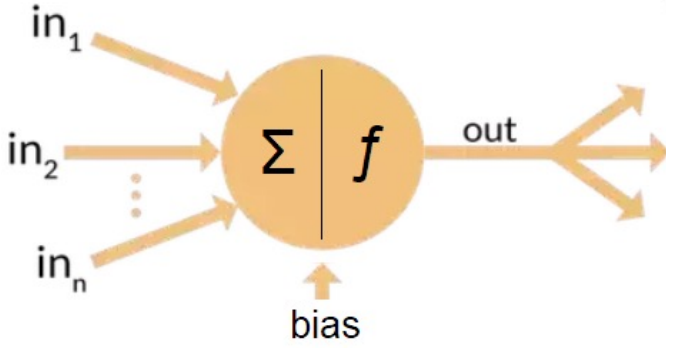


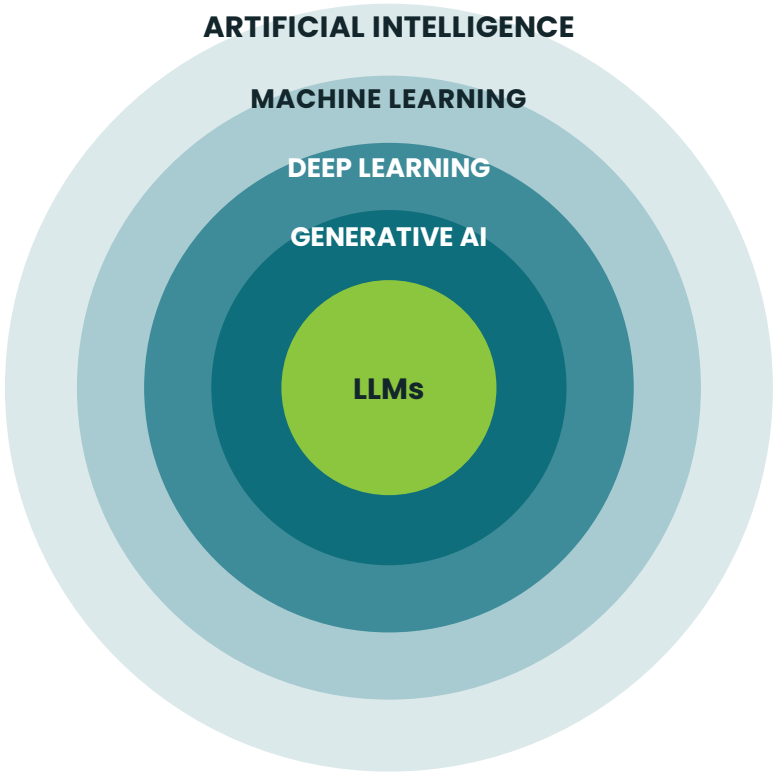
Frank Rosenblatt works on the "electronic profile analyzing computer" – a precursor to the perceptron.

Cornell chronical: Division of Rare and Manuscript Collections

What is AI?







- **AI**

  Systems that make predictions and do tasks that usually need human intelligence.

- **Machine learning (ML)**

  Learns from data to do tasks without being explicitly programmed.

- **Deep learning (DL)**

  Neural networks that recognise complex patterns, the base of modern AI.

- **Generative AI**

  Creates new content (text, images, audio) from patterns it has learned.

- **Large language models**

  GenAI trained on massive text; general-purpose systems (Claude, ChatGPT, Gemini, Alfafold,).

Artificial Neural Network (ANN) → Transformer architecture models

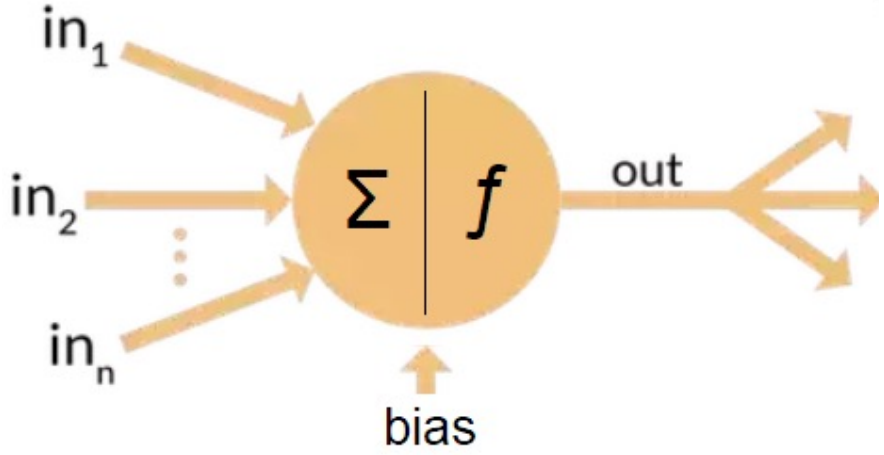
From broad to specific: AI → machine learning → deep learning → generative AI → LLMs.

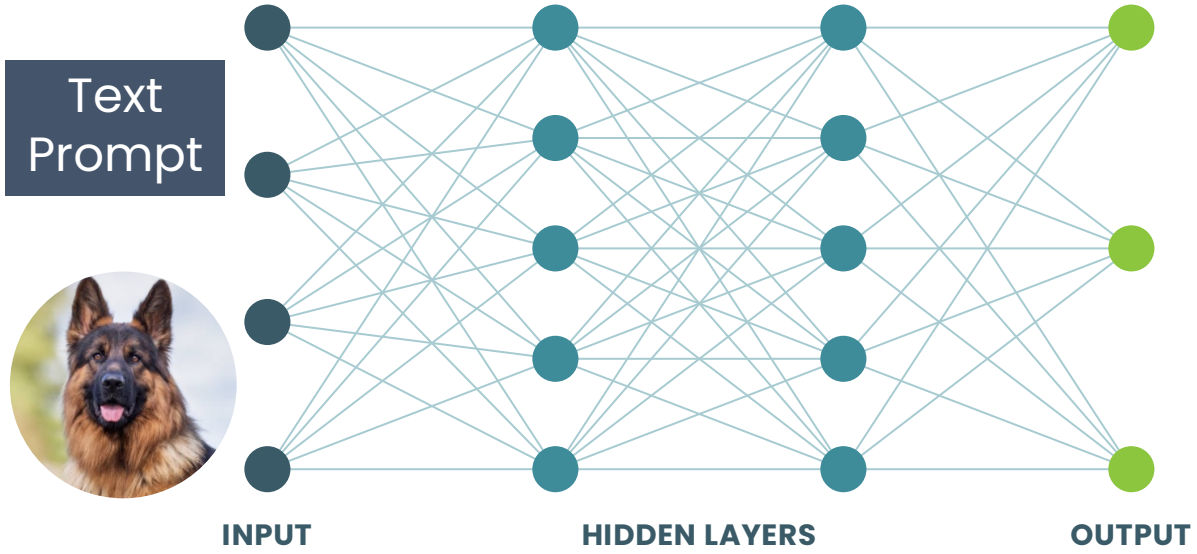
Standard AI taxonomy; nested AI ⊃ ML ⊃ DL hierarchy.

UNDER THE HOOD

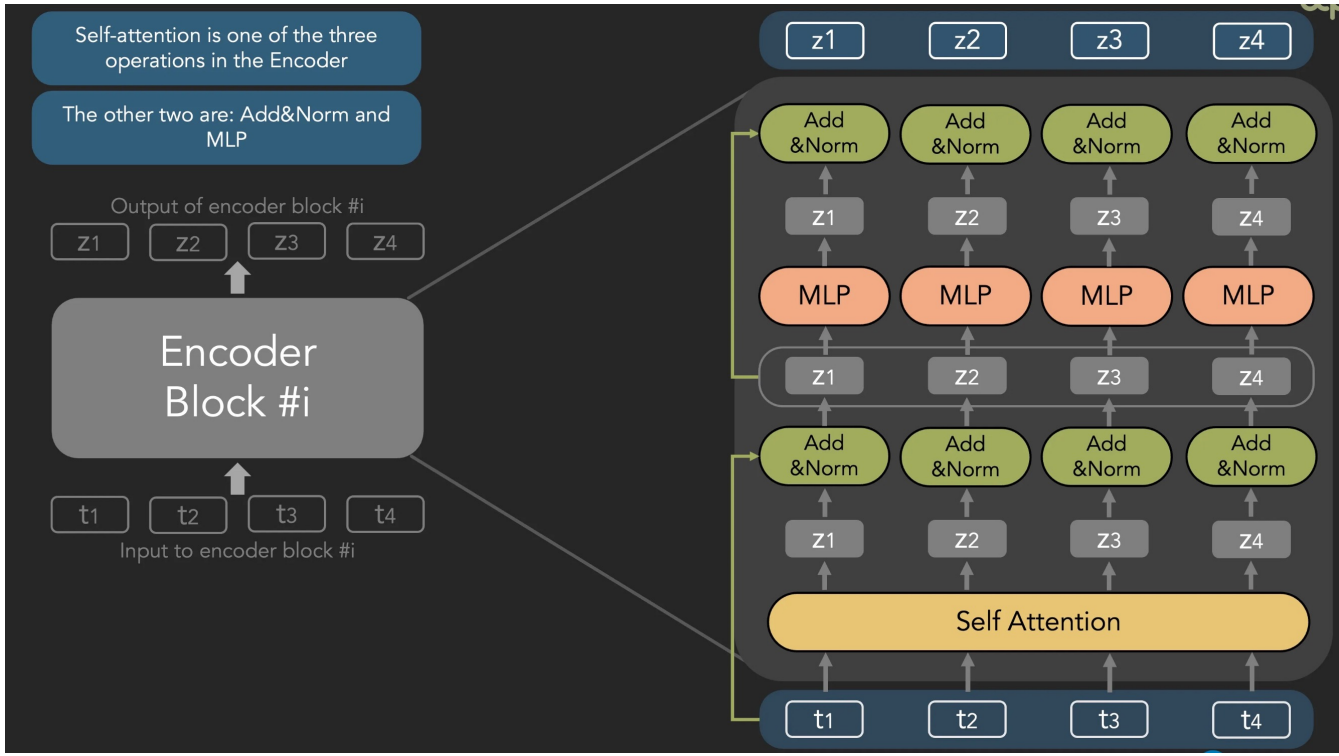
Inside a neural network

Layers of simple units, tuned by data, turn inputs into predictions.





A TRANSFORMER BLOCK



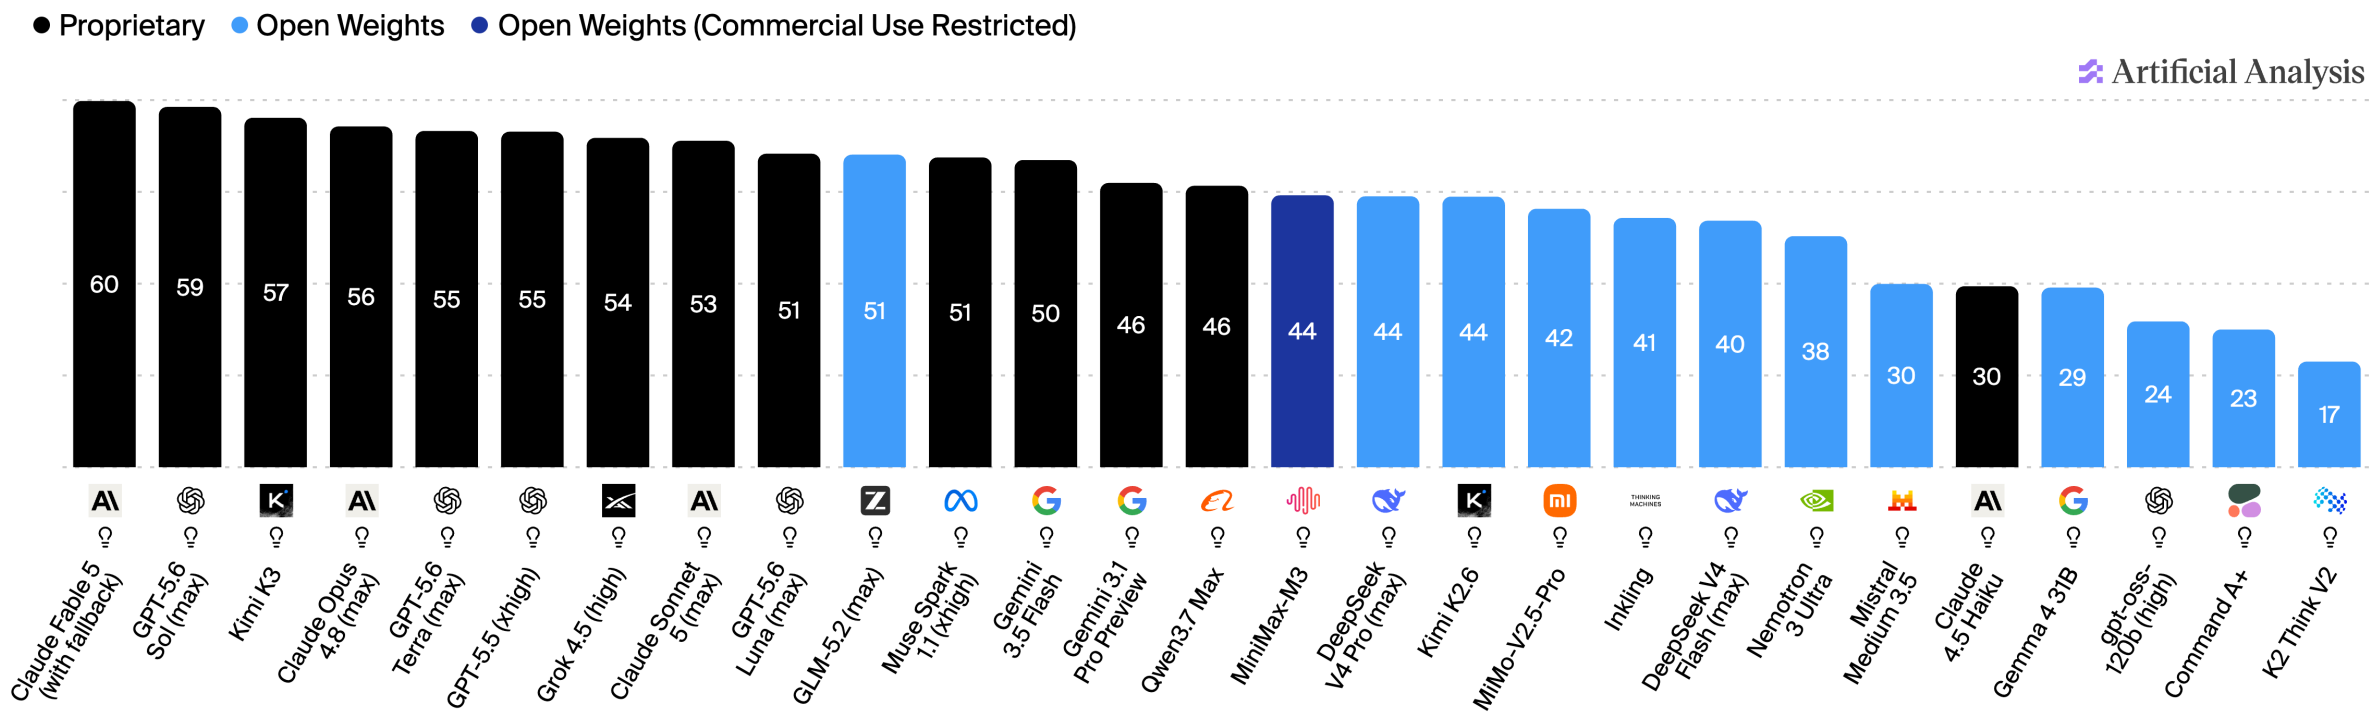
Every connection has a weight; training nudges millions of them until predictions match reality.

Neural networks: Rumelhart et al. (1986). Transformer block: Vaswani et al. (2017).

Model releases (Proprietary vs Open)

Artificial Analysis Intelligence Index by Open Weights / Proprietary

Artificial Analysis Intelligence Index v4.1 incorporates 9 evaluations: GDPval-AA v2, τ^3-Banking, Terminal-Bench v2.1, SciCode, Humanity's Last Exam, GPQA Diamond, CritPt, AA-Omniscience, AA-LCR

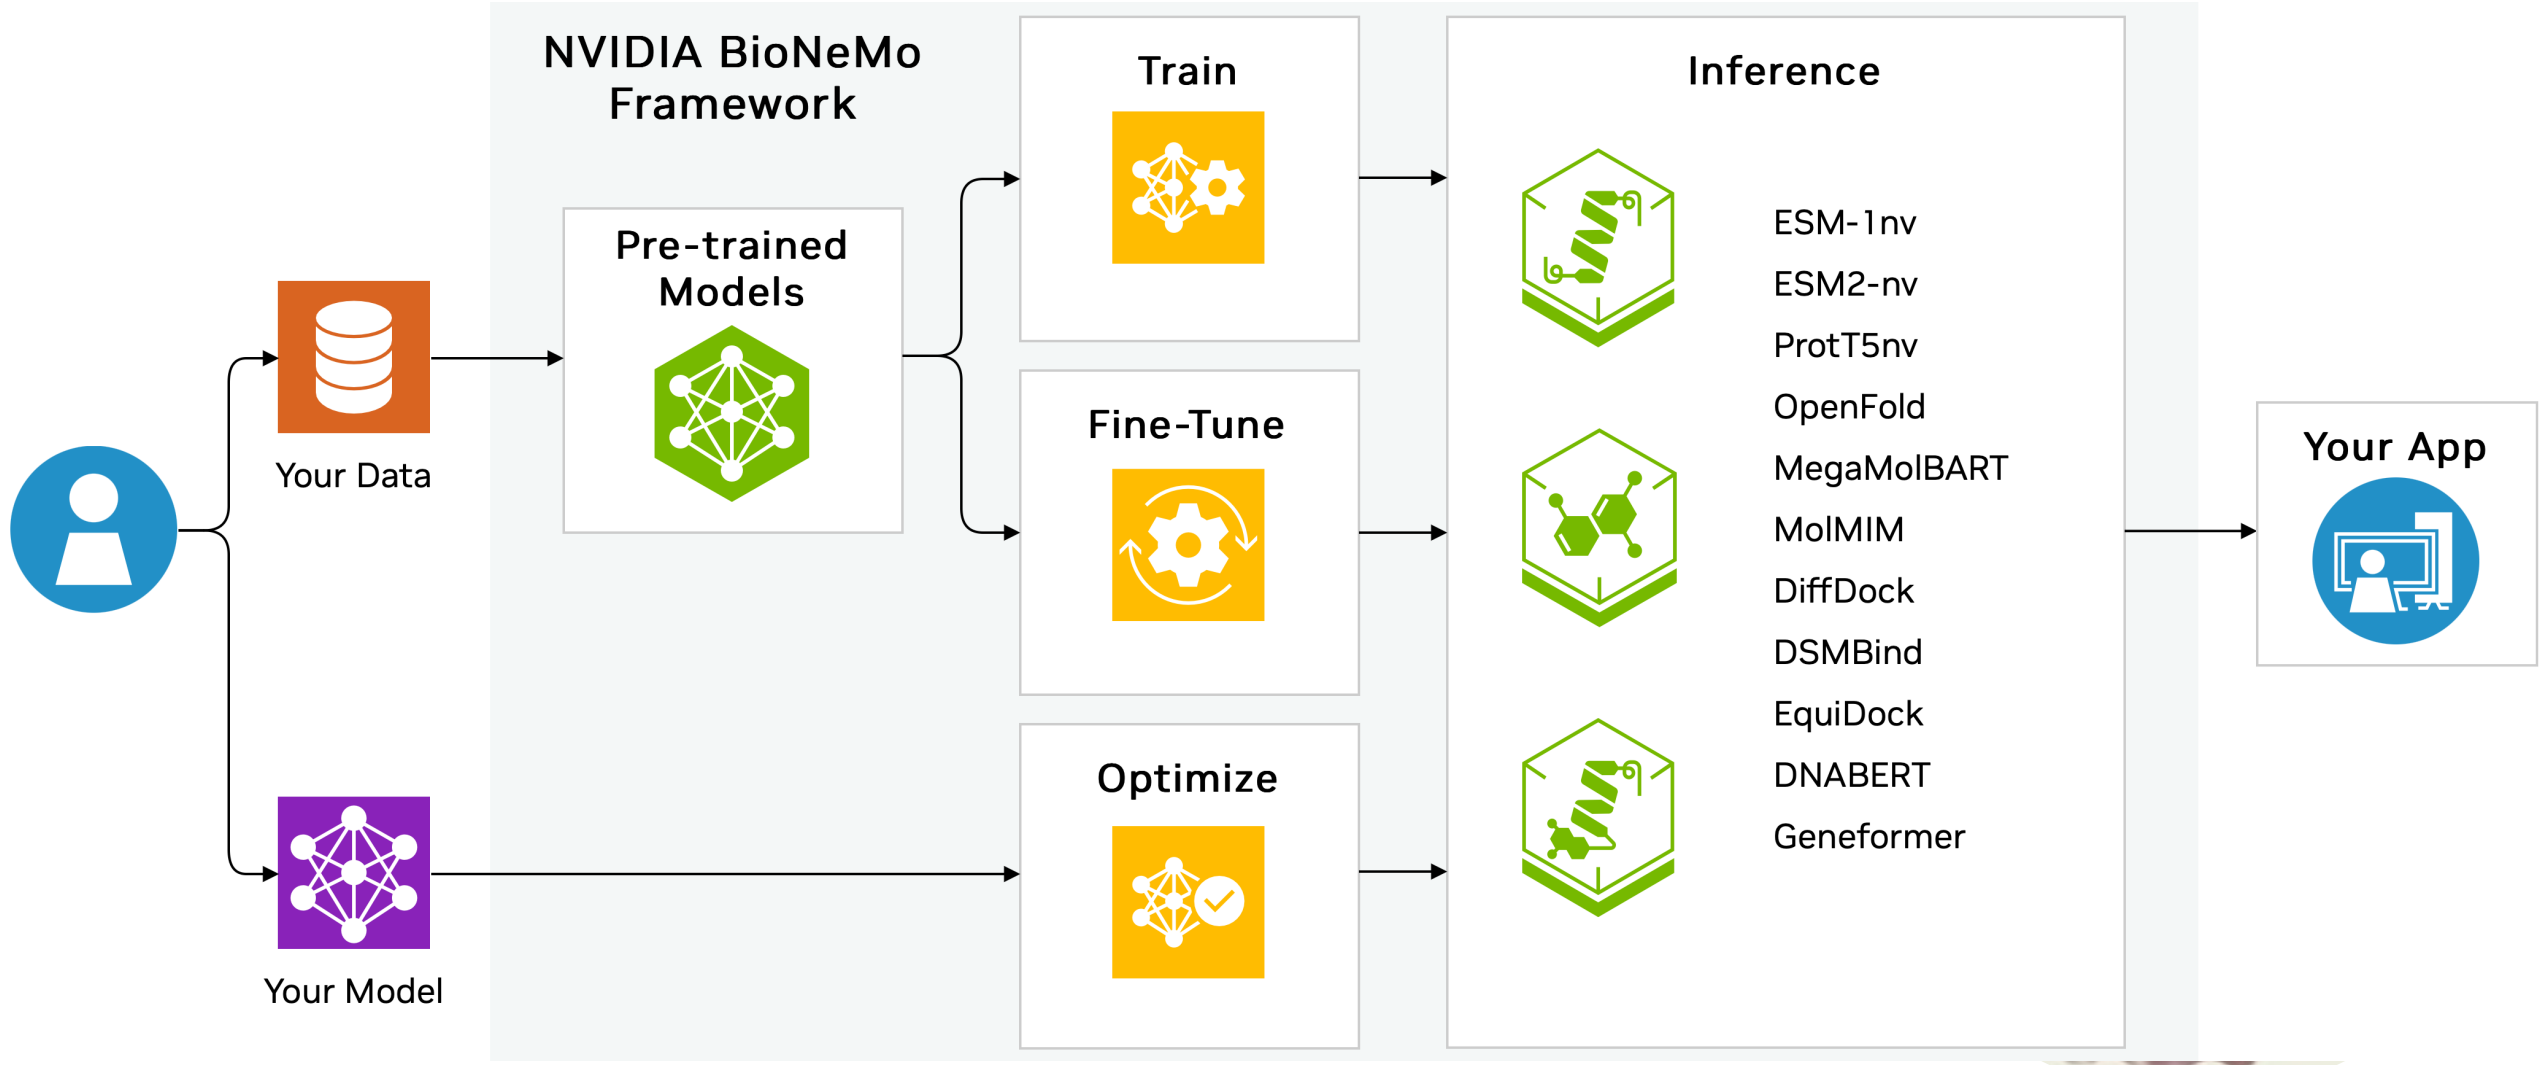


Reasoning models are indicated by a lightbulb icon

The frontier moves fast, and open models, only months behind, keep the options open.

Stanford AI Index 2025; Epoch AI. Counts of notable foundation models per year (approximate).

Genomic models with varying DL architectural frameworks



Sir Demis Hassabis, Google DeepMind

Open-weight vs proprietary

OPEN-WEIGHT · SELF-HOSTABLE

Qwen 3.5 · DeepSeek V3.2 / R1 · Llama 4 · Mistral · Gemma 3

- Runs on your own machines, so data never leaves.
- Free to license (Apache-2.0, MIT); fine-tuned locally for specific institutional tasks.
- **Small models (Gemma 3) run on modest hardware, even offline.**
- Best for sensitive data, sovereignty and local languages.

PROPRIETARY · API / CLOUD

Claude Opus 4.8 · GPT-5.5 · Gemini 3.1 Pro

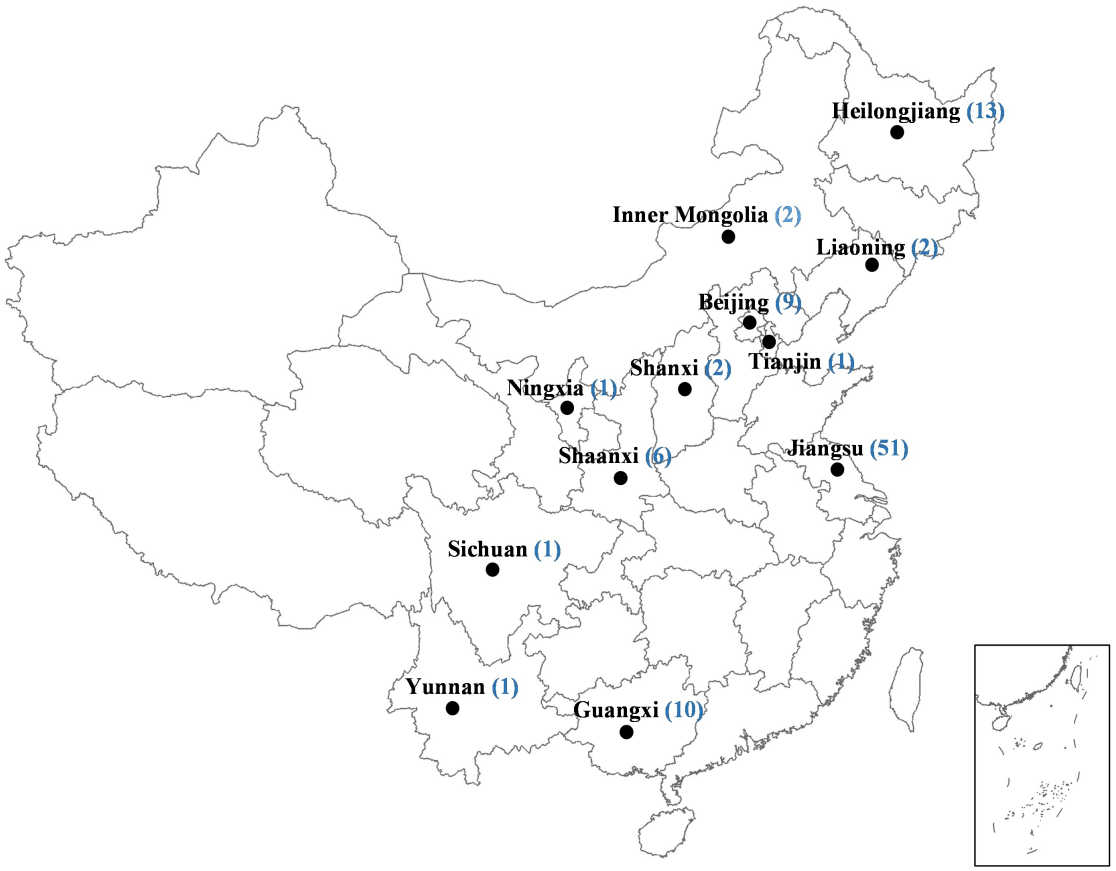
- Highest capability, nothing to install or maintain.
- Pay per use; prices fell 30–60% over the past year.
- **Data leaves your premises, so check policy and consent.**
- Best for peak reasoning, drafting and one-off analysis.

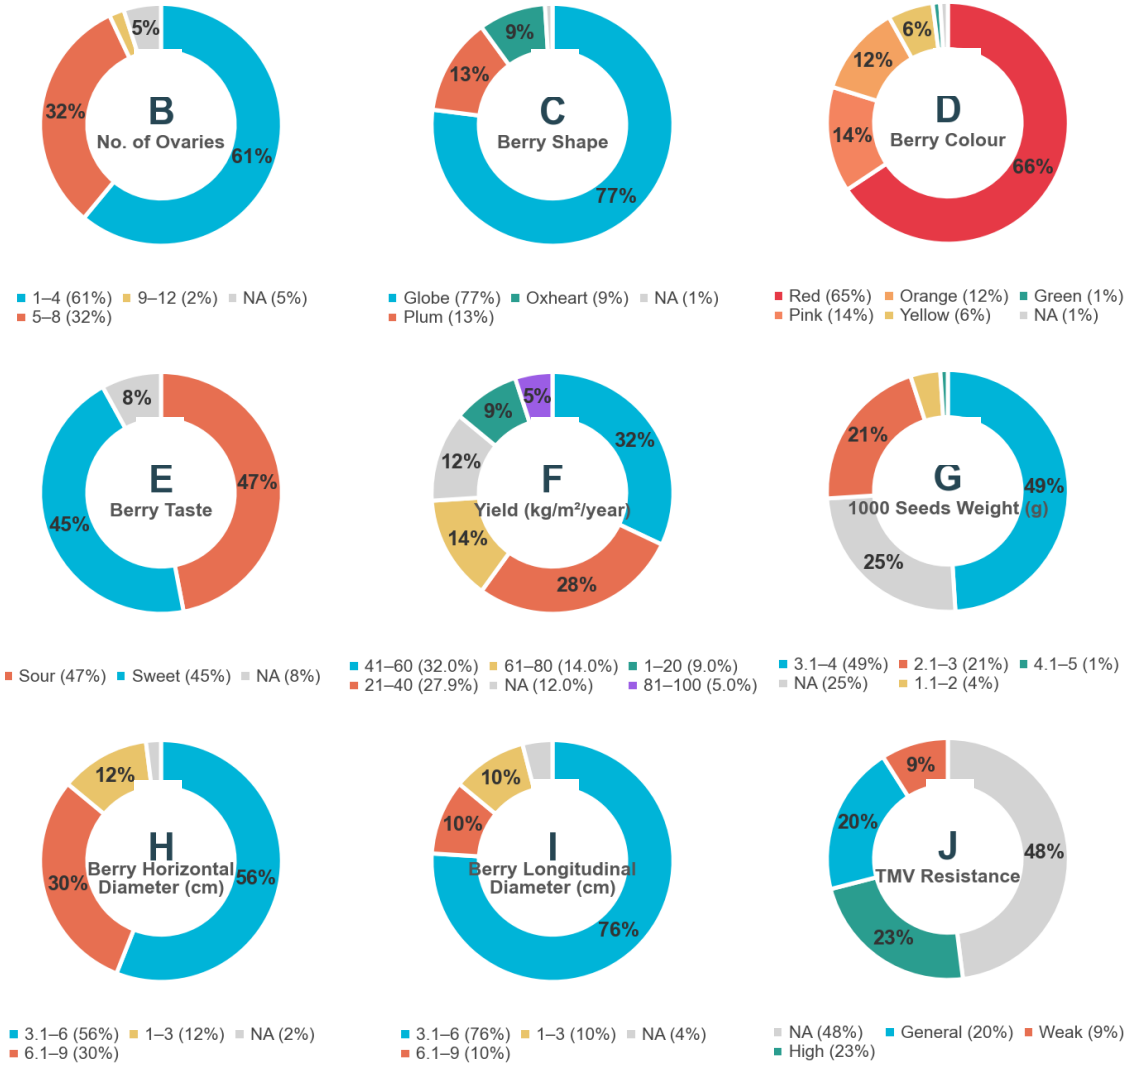
Rule of thumb: sensitive or routine → open and local; hardest problems → proprietary, with care.

Model capabilities per public benchmarks, mid-2026 (Artificial Analysis; Hugging Face Open LLM leaderboards).

DL & ML use case: Factors that shape the seed microbiome?

- *100 Genotypes, 12 provinces (China)*
- *16S rRNA amplicon sequencing of seed endosphere*

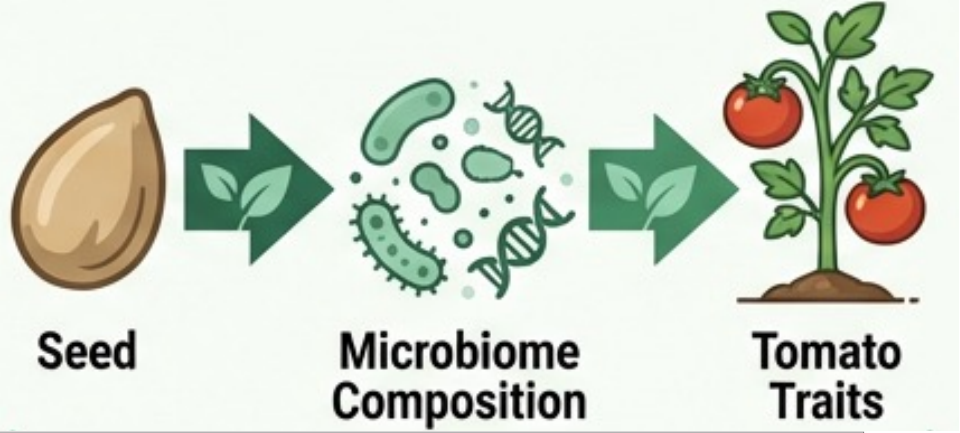


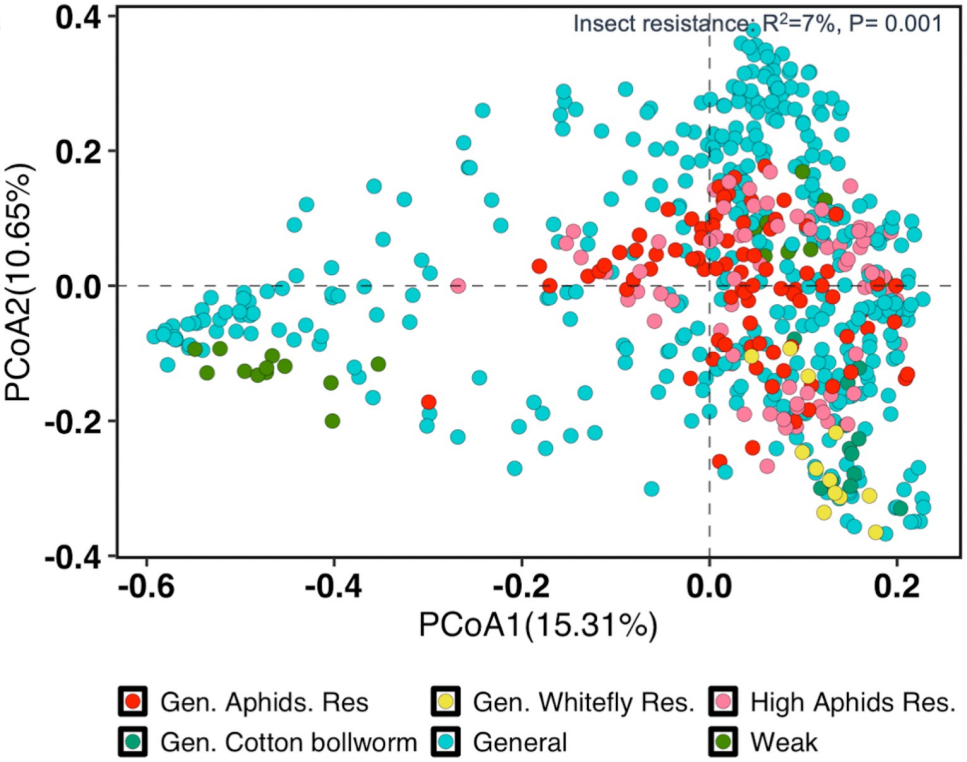


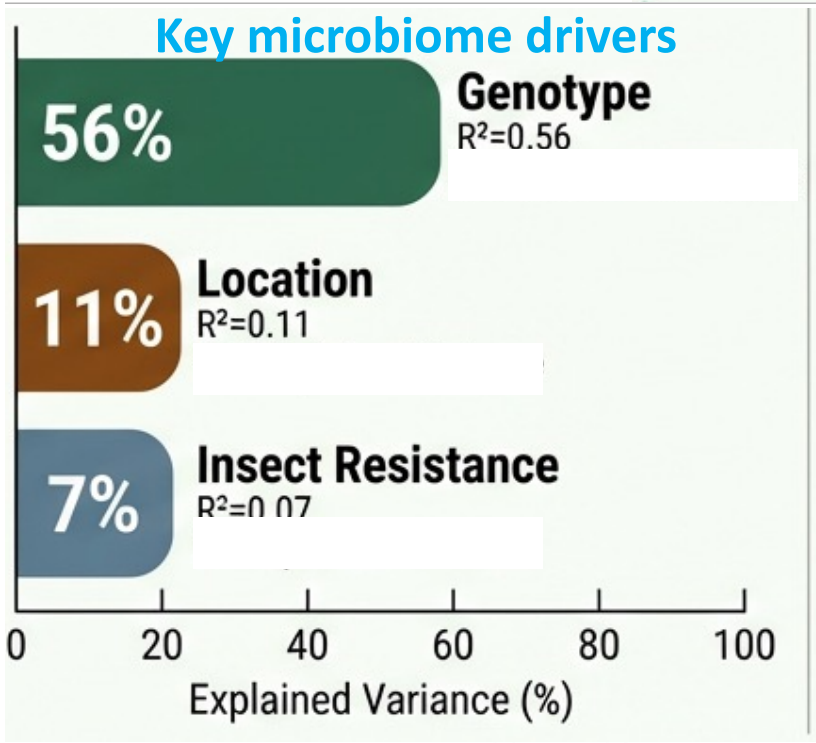
Chen, Olimi, Cernava et al. 2025

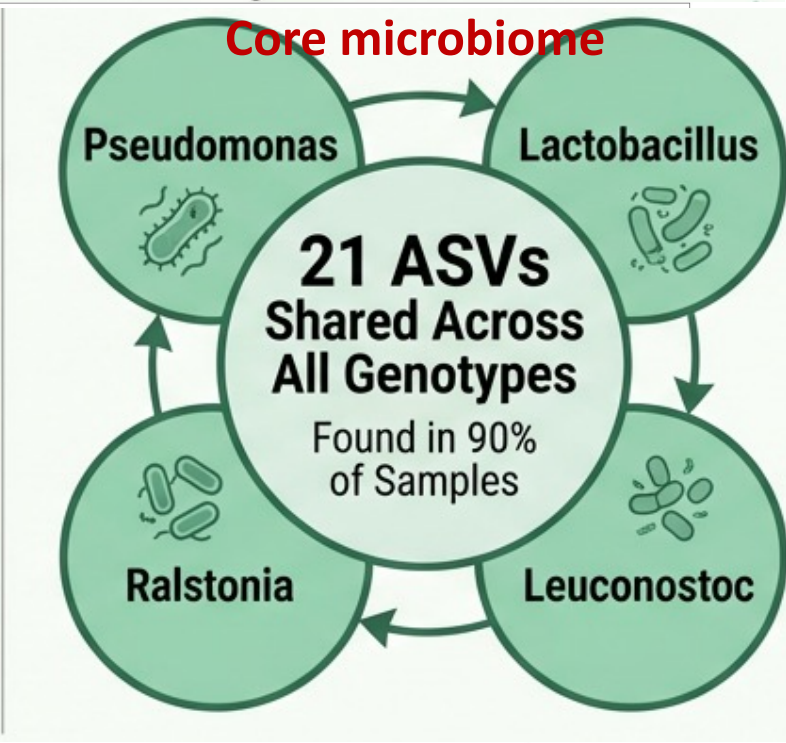
Examined how **host genetics & environment**, as well as host traits shape the seed microbiome

ML classification tasks



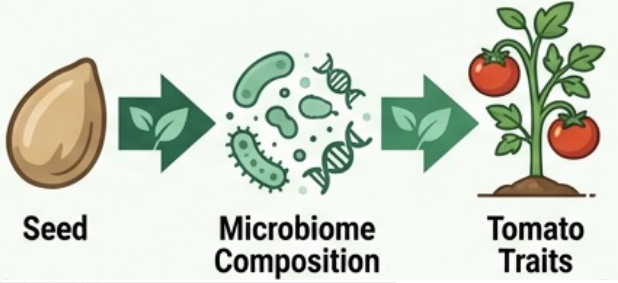


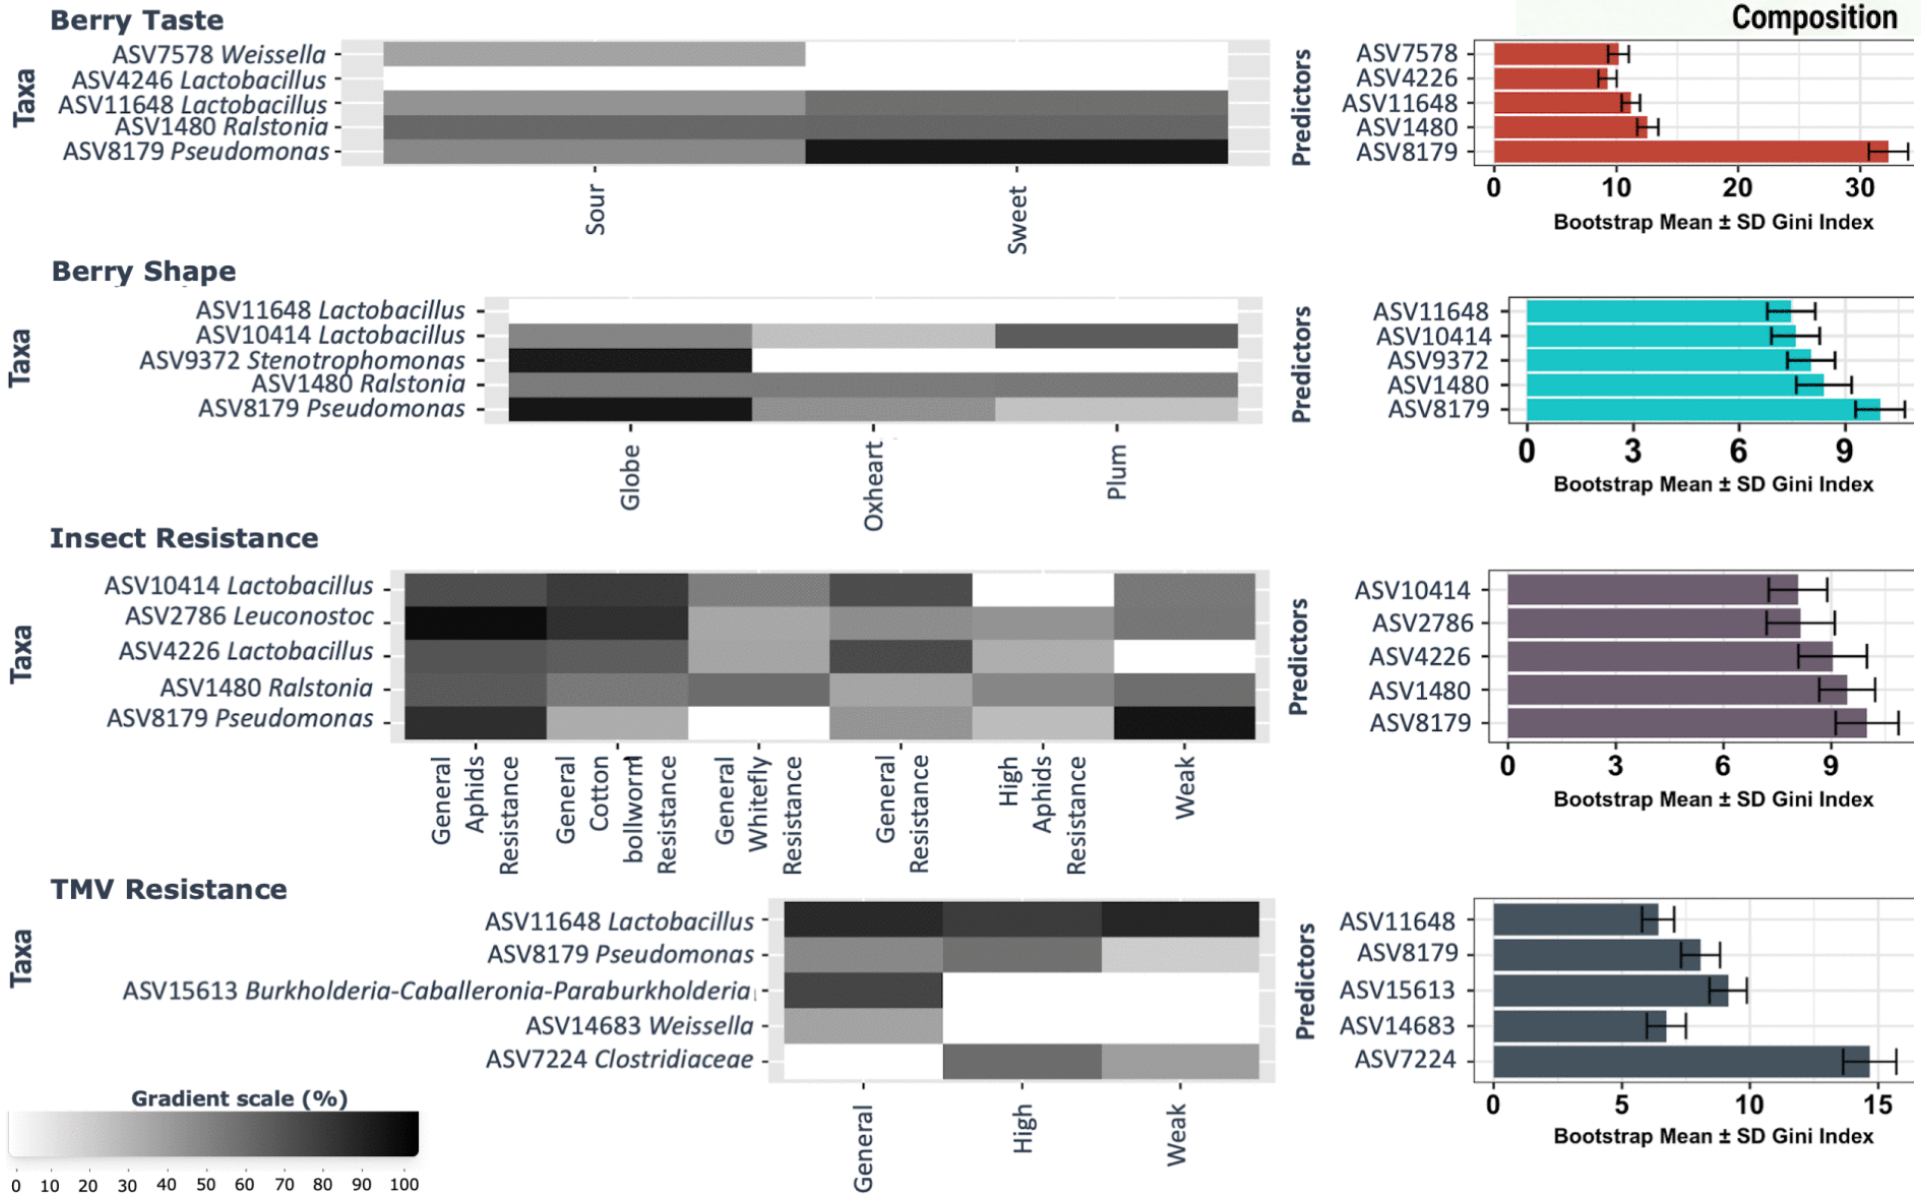




Tomato seeds host a high bacterial diversity, whose structure is mainly shaped by host genotype

ML predictive task: biomarker discovery





Chen, Olimi, Cernava et al. 2025

Agents: Models + Tools capabilities



Surveillance

flags outbreaks · drafts alerts

Breeding

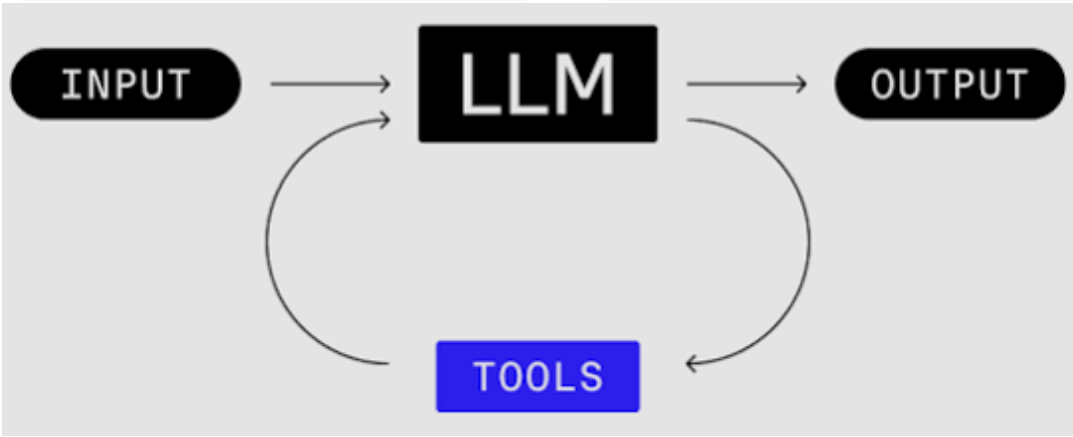
ranks crosses · drafts reports

Lab

runs QC · flags anomalies

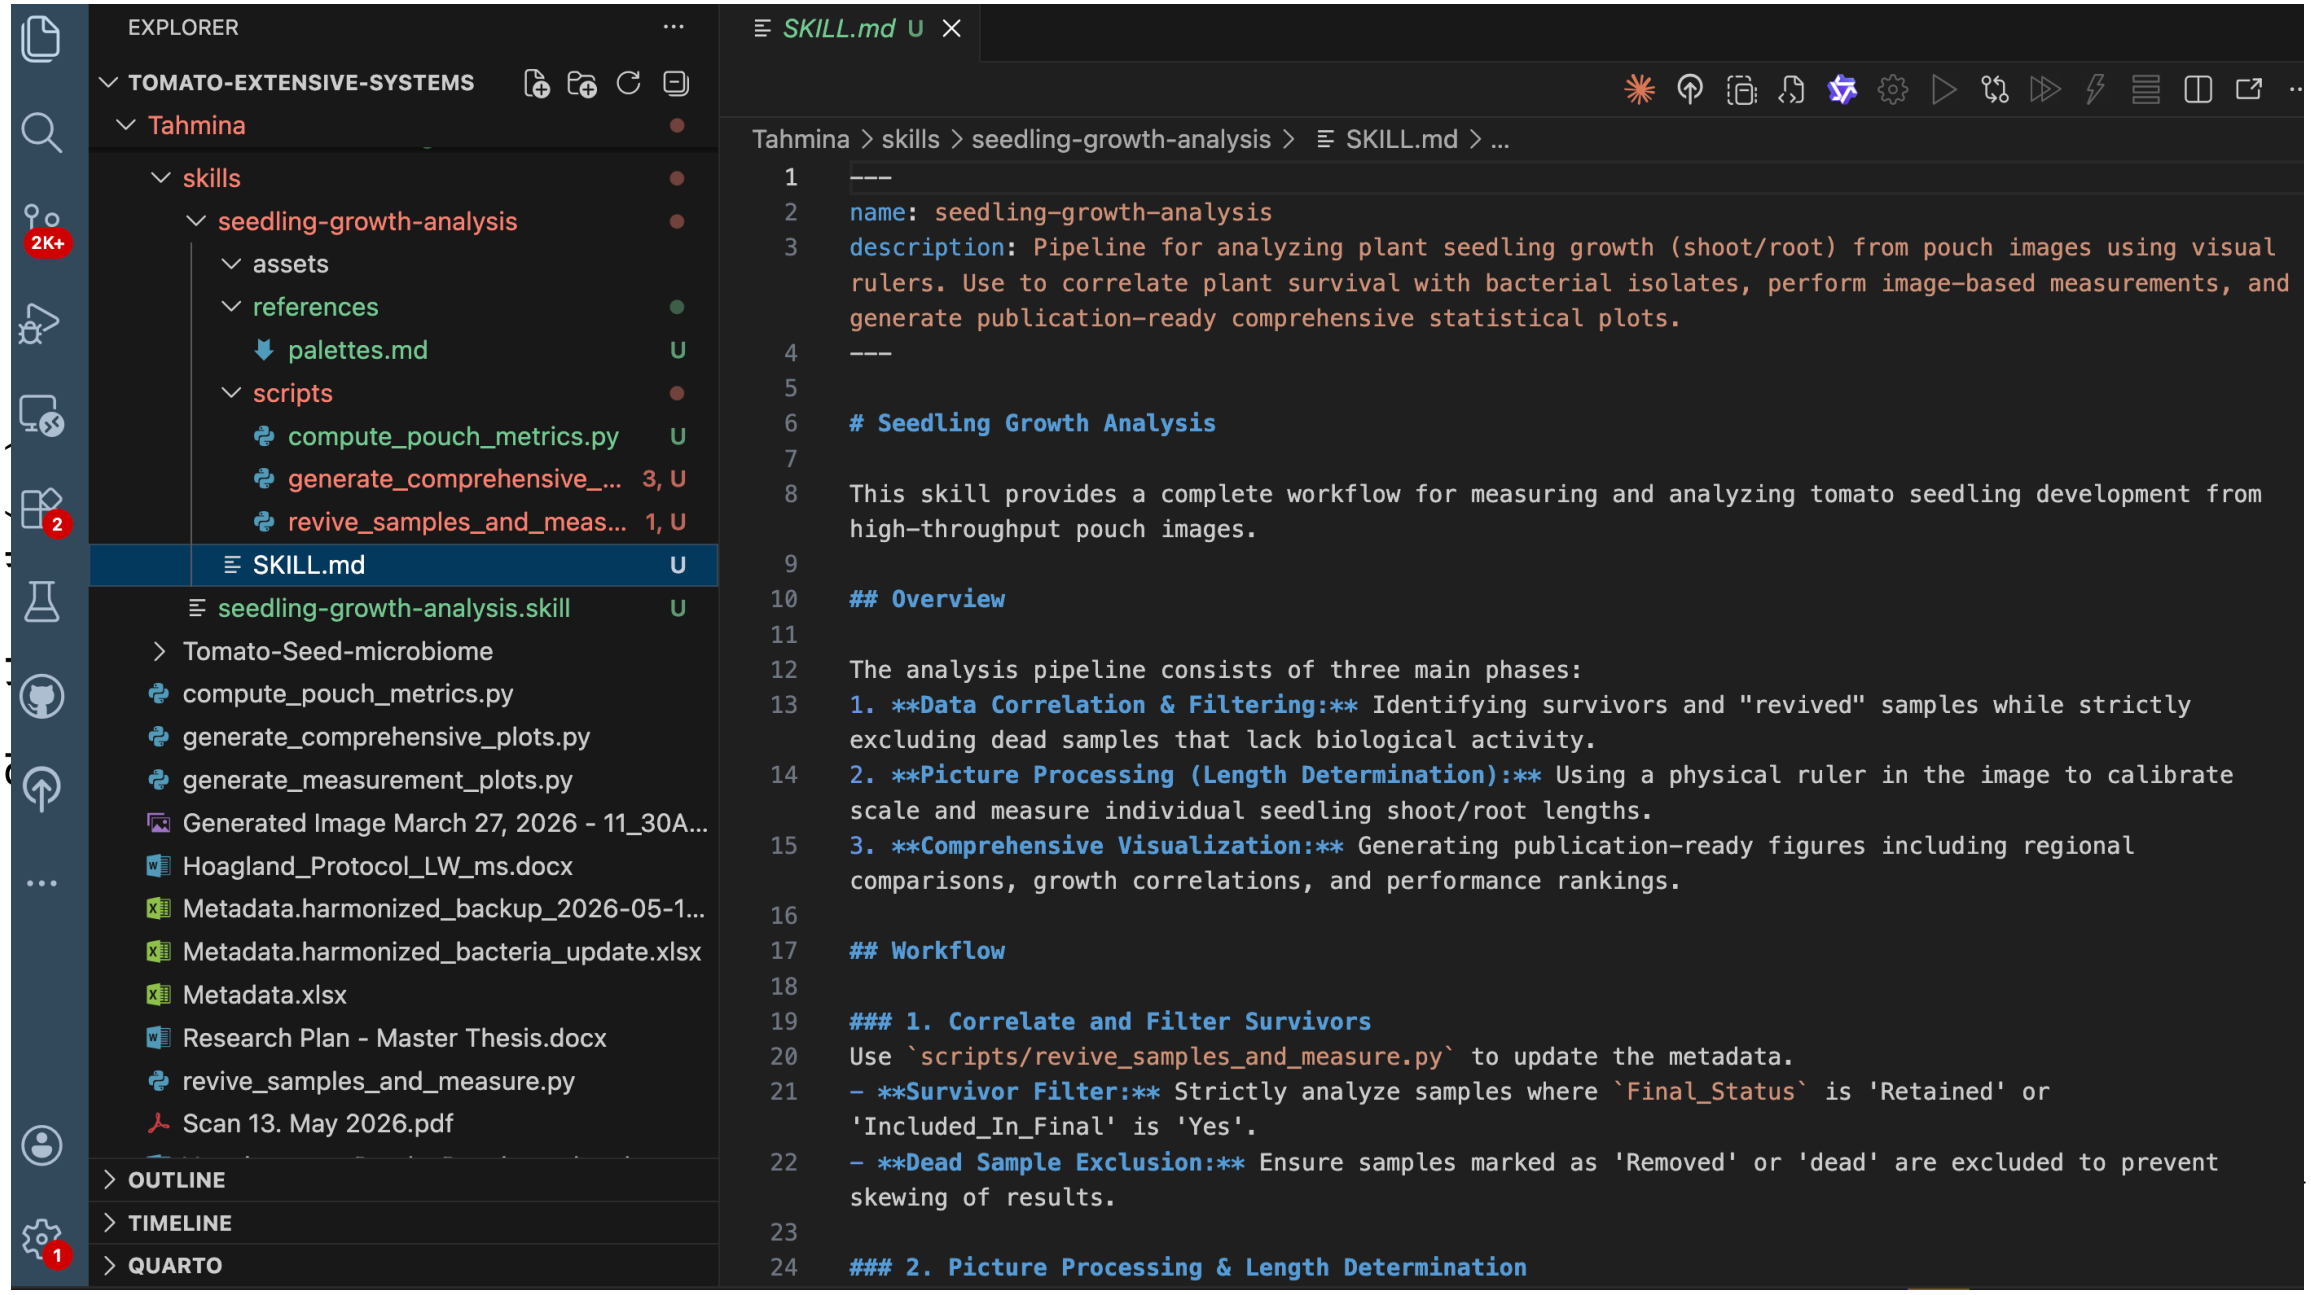
Knowledge

answers with citations

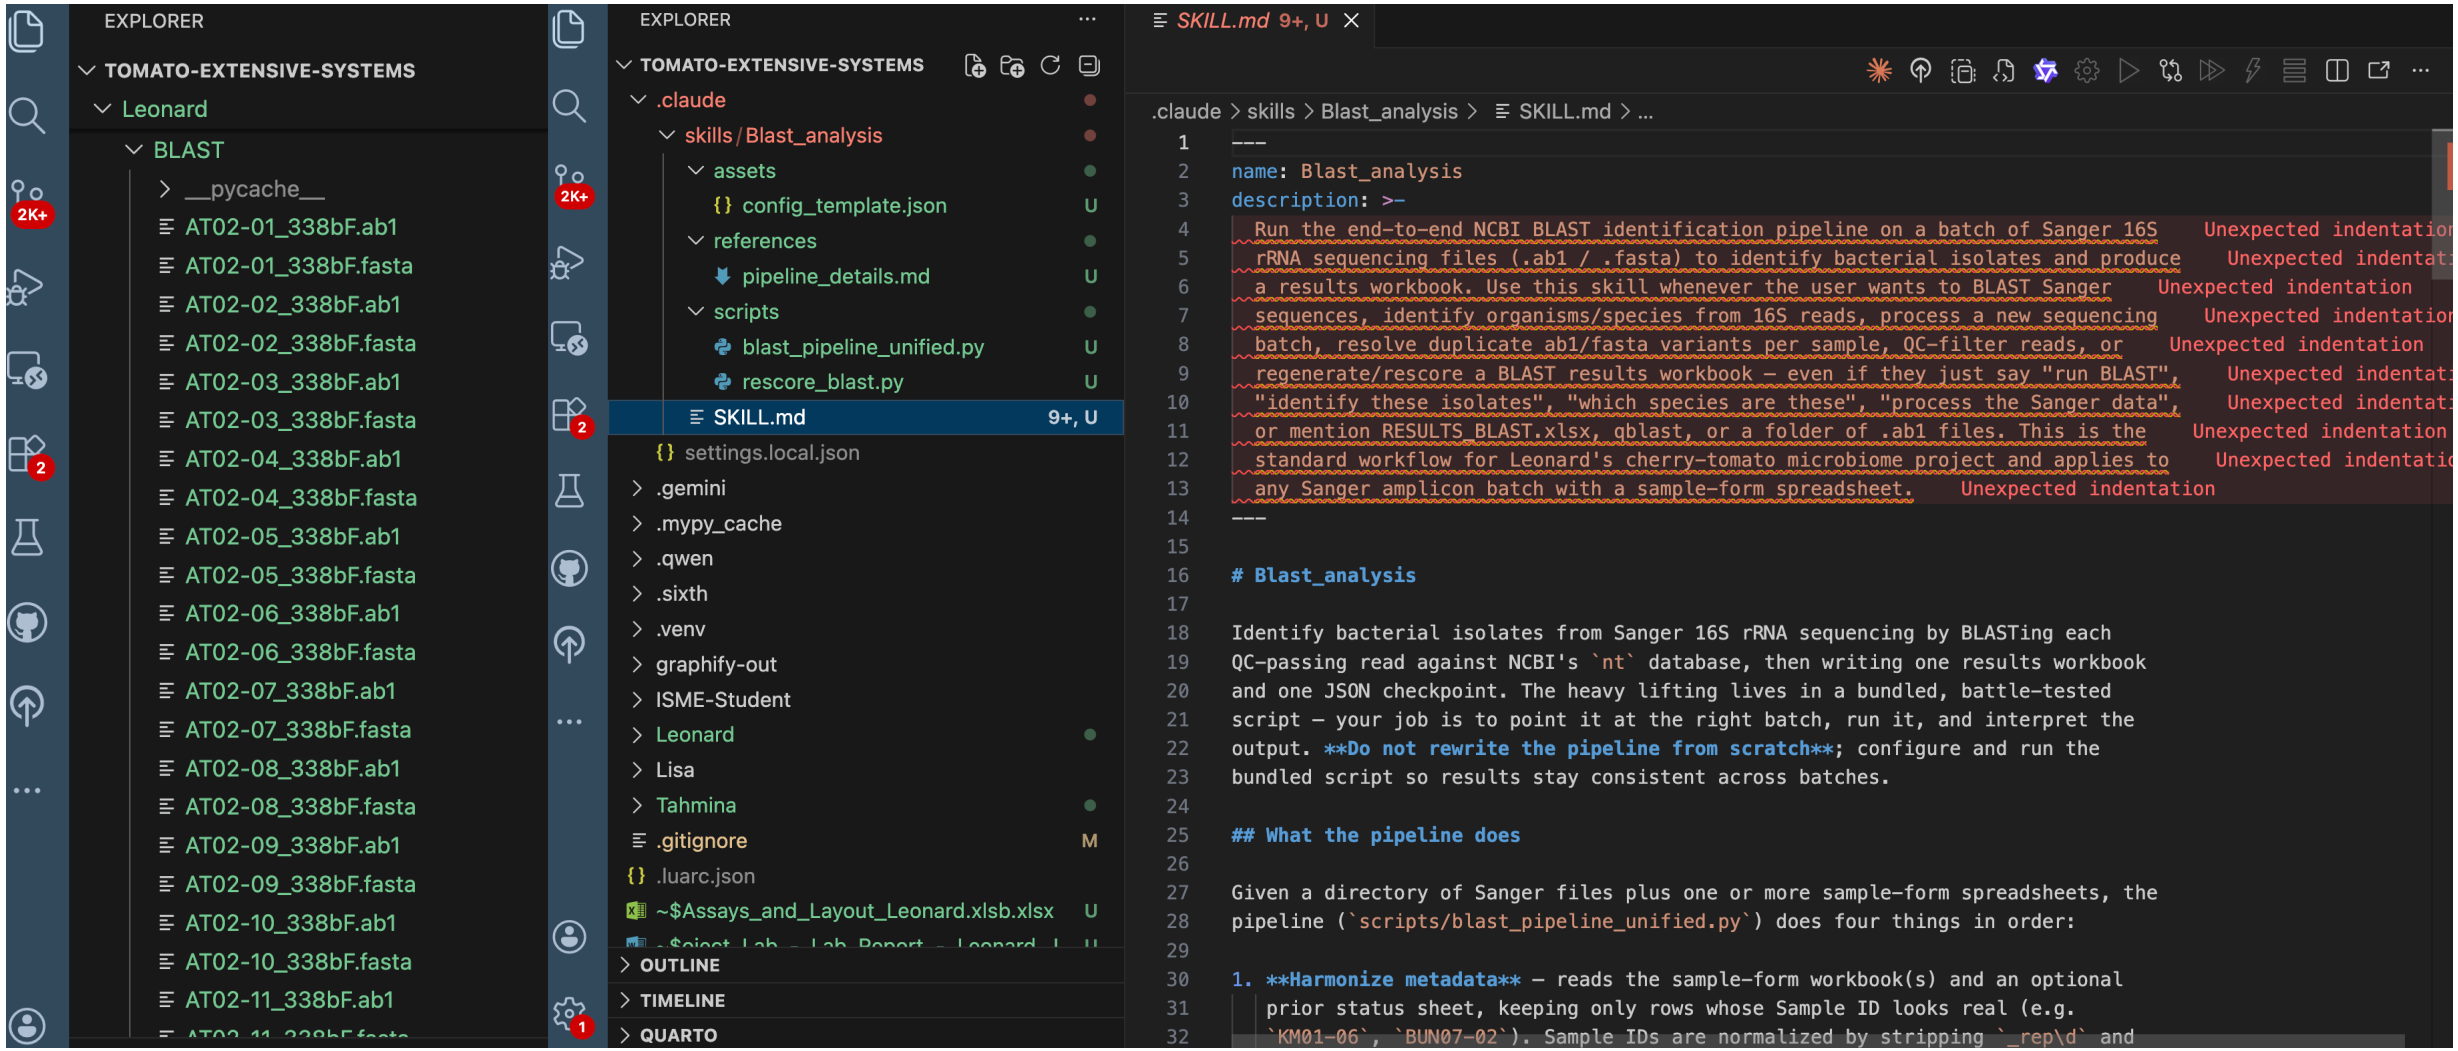


An agent plans, uses tools and carries a task through (Human-in-loop-control)

Agentic application in image analysis (Claude code)



Skills help to execute repetitive tasks



Extrapolation to: Analyse genomes, Project management, Data analysis, Research design, etc.,

Claude Science, an AI workbench for scientists

IUCN red list analysis by taxonomy



When writing this paper that was published at Nature Revie
earlier draft, I wrote something to the effect of species on th
"downgraded" to less threatened categories (e.g., critically e
endangered to vulnerable). An astute reviewer picked up on
it — if I'm going to say such a thing, I'd need to back that up
analysis.



Genome engineering in biodiversity conserva

STEPHEN D. TURNER • 18 JULY 2025

Read full story →

I never had the time to do this analysis, so I just deleted the
paper without any analysis. But I always thought it would ha
particular taxa "downgraded" more than other taxa? How do
and geography? How might those downgradings correlate o
of genetic diversity, as we discussed in the paper?

I gave it some rough guidance:

> I'd like to pull some data from the IUCN API and do som
> trends in categorization by taxonomy class, preferably usi

AI-Biosecurity Literature Meta-analysis | **IUCN API Data Analysis by Taxonomy**

Your last message

Excellent — this gives me what I need. The clas ssessments (including historical ones, flagged by `latest`), with `year_published`, `red_list_category_code`, and a `scopes` field (Global vs regional). That's ideal for framing 3: `latest=TRUE` + Global scope gives the current snapshot, and the full set of `year_published` × category gives the temporal signal. Let me check the scope structure and gauge the data volume.

✓ Checking scope structure and Mammalia volume — 16 lines of output

Now I understand the data model well. Key findings:

- The per-class endpoint returns **all** assessments (historical + current, global + regional), 100 per page. `latest=TRUE` + Global scope = current snapshot; full set + `year_published` = temporal trends.
- Historical assessments carry **legacy category codes** (`E`, `V`, `LR/lc`, `LR/nt`, etc.) alongside modern ones (`CR`, `EN`, `VU`, `NT`, `LC`, `DD`) — I'll need to handle both.
- Volume is substantial: Mammalia alone has >20,000 assessment records (counting regional + historical). The full database is 172,620 assessments ≈ ~1,730 pages.

Let me check timing and rate-limit behavior with a quick sequential burst before committing to the full pull strategy.

Plan ready for your review

Pull IUCN Red List data and analyze extinction-risk categorization trends by taxonomic class (R)

Open | Approve

high confidence

Describe changes to the plan…

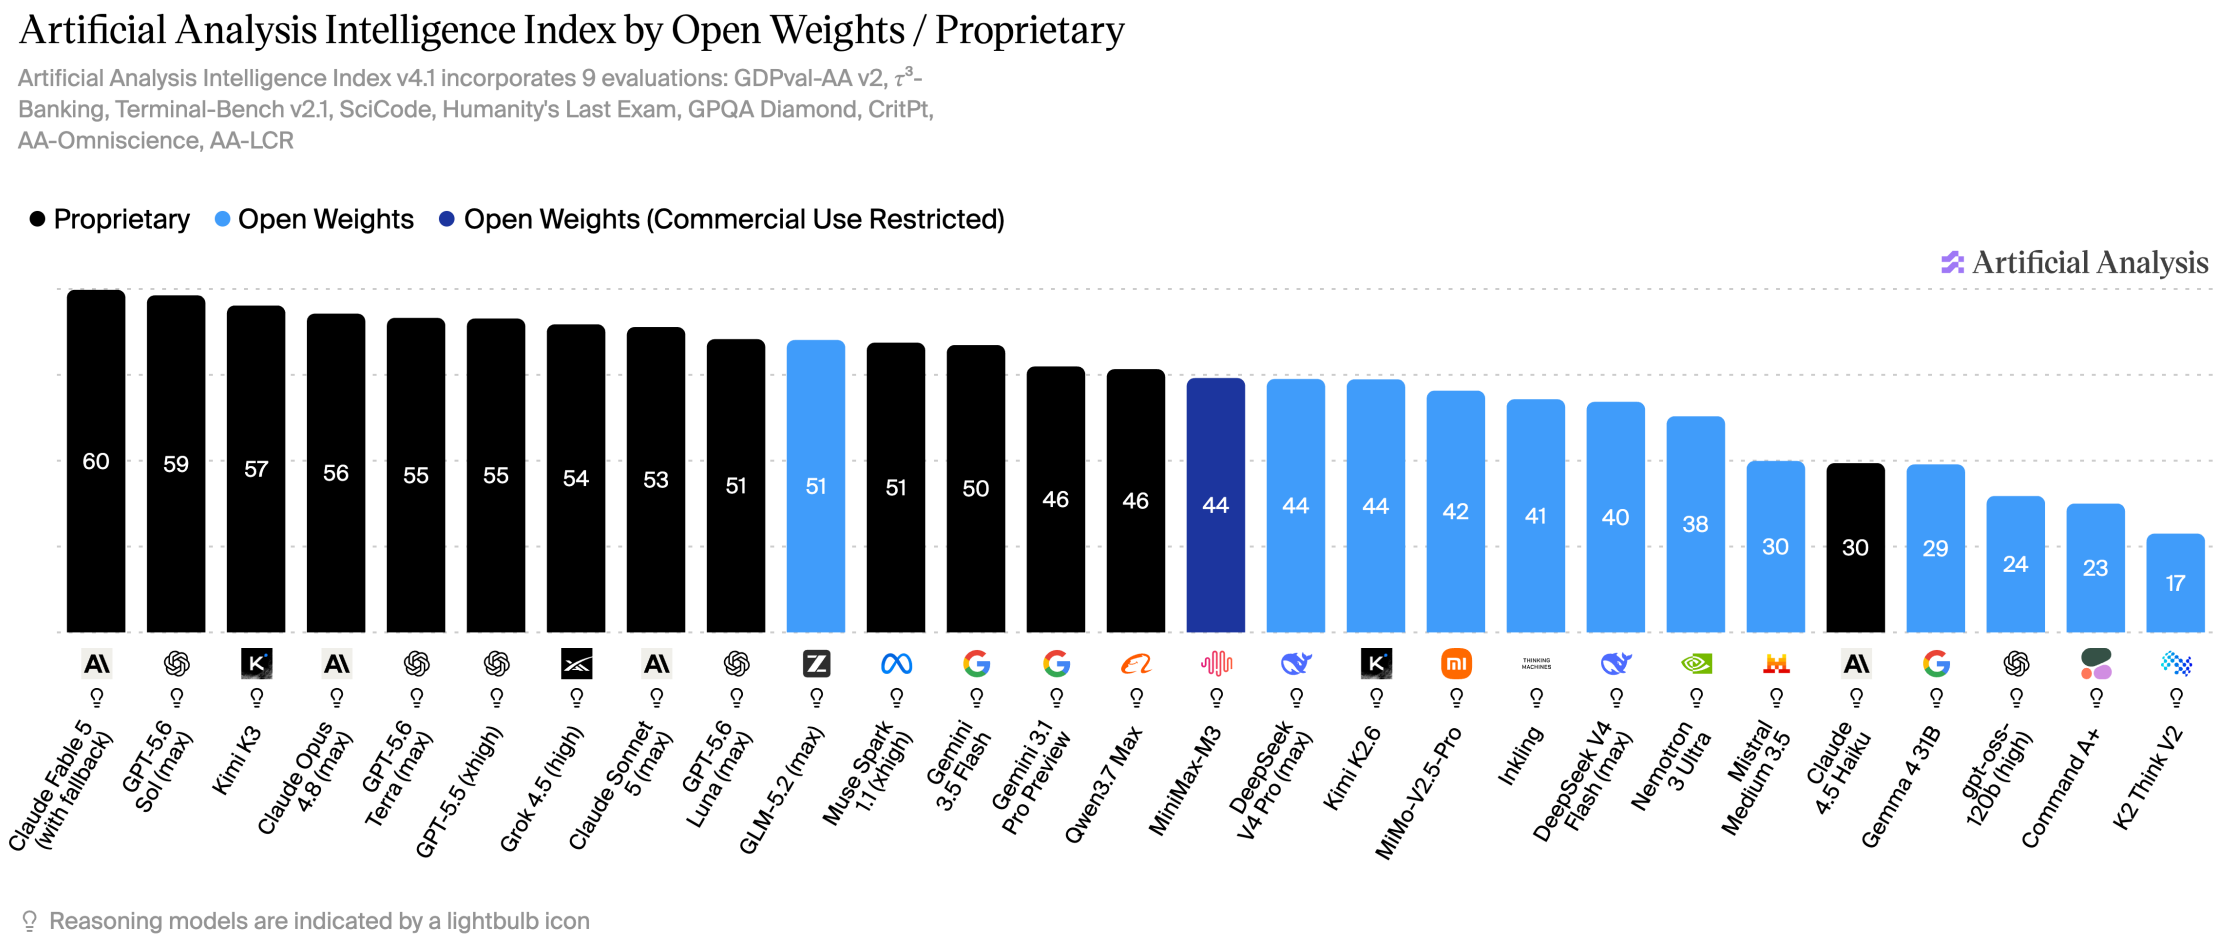
Notebook | plan_pull-iucn-red-list-data-and-ana…

R | Agent SDK

[7] r

```r
# Try common filter params to avoid pulling all historical+regional records
tries <- list(
  list(latest="true"),
  list(scope_code="1"),
  list(scopes="1"),
  list(year_published="2024")
)
for (q in tries) {
  q$page <- 1
  rr <- GET(paste0(base, "/taxa/class/MAMMALIA"), hdr, query=q)
  oo <- tryCatch(fromJSON(content(rr,"text",encoding="UTF-8"), simplifyVector=FALSE), error=function(e) NULL)
  n <- if(!is.null(oo$assessments)) length(oo$assessments) else NA
  lat <- if(!is.null(oo$assessments)) sum(sapply(oo$assessments, function(a) isTRUE(a$latest))) else NA
  cat(sprintf("query=%s -> status %d, n=%s, latest_in_page=%s\n",
              paste(names(q),unlist(q),sep="=",collapse="&"), status_code(rr), n, lat))
}
```

Leading frontier models in agentic workflows (i.e., claude code/cowork, codex, gemini, copilot…)



Reasoning models are indicated by a lightbulb icon

For most research work, an open model can give ~90% of the capability at a tiny fraction of the cost.

Intelligence Index vs published API price, mid-2026 (approximate; open models also self-hostable).

Why this matters for Uganda & NARO?



>70% of Ugandans work in agriculture. Anchors food security & the economy

<5%

Global crop-AI training data is African. **Models don't see data (genomic: plant and animal genetic resources..)**

2025

MAAIF opened its National Digital Agriculture Strategy consultations.

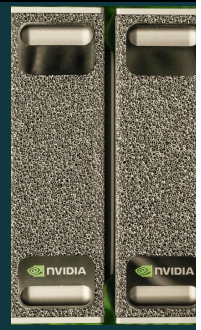
AI is a research multiplier. NARO should strategically invest and position to harness the benefits that the technology offers.

THE ENABLER

Hardw

The frontier

NVIDIA Asia Paci
One DGX Spark o
Four is a local AI f



BUSINESS INSIDER AFRICA

MARKETS LEADERS CAREERS LIFESTYLE

The swift global adoption of AI, alongside Africa's rapid tech growth, increasingly suggests the feasibility of an AI factory on the continent.



Another 'first AI factory in Africa'? Uganda is said to have made its entry into the AI race

cally
els

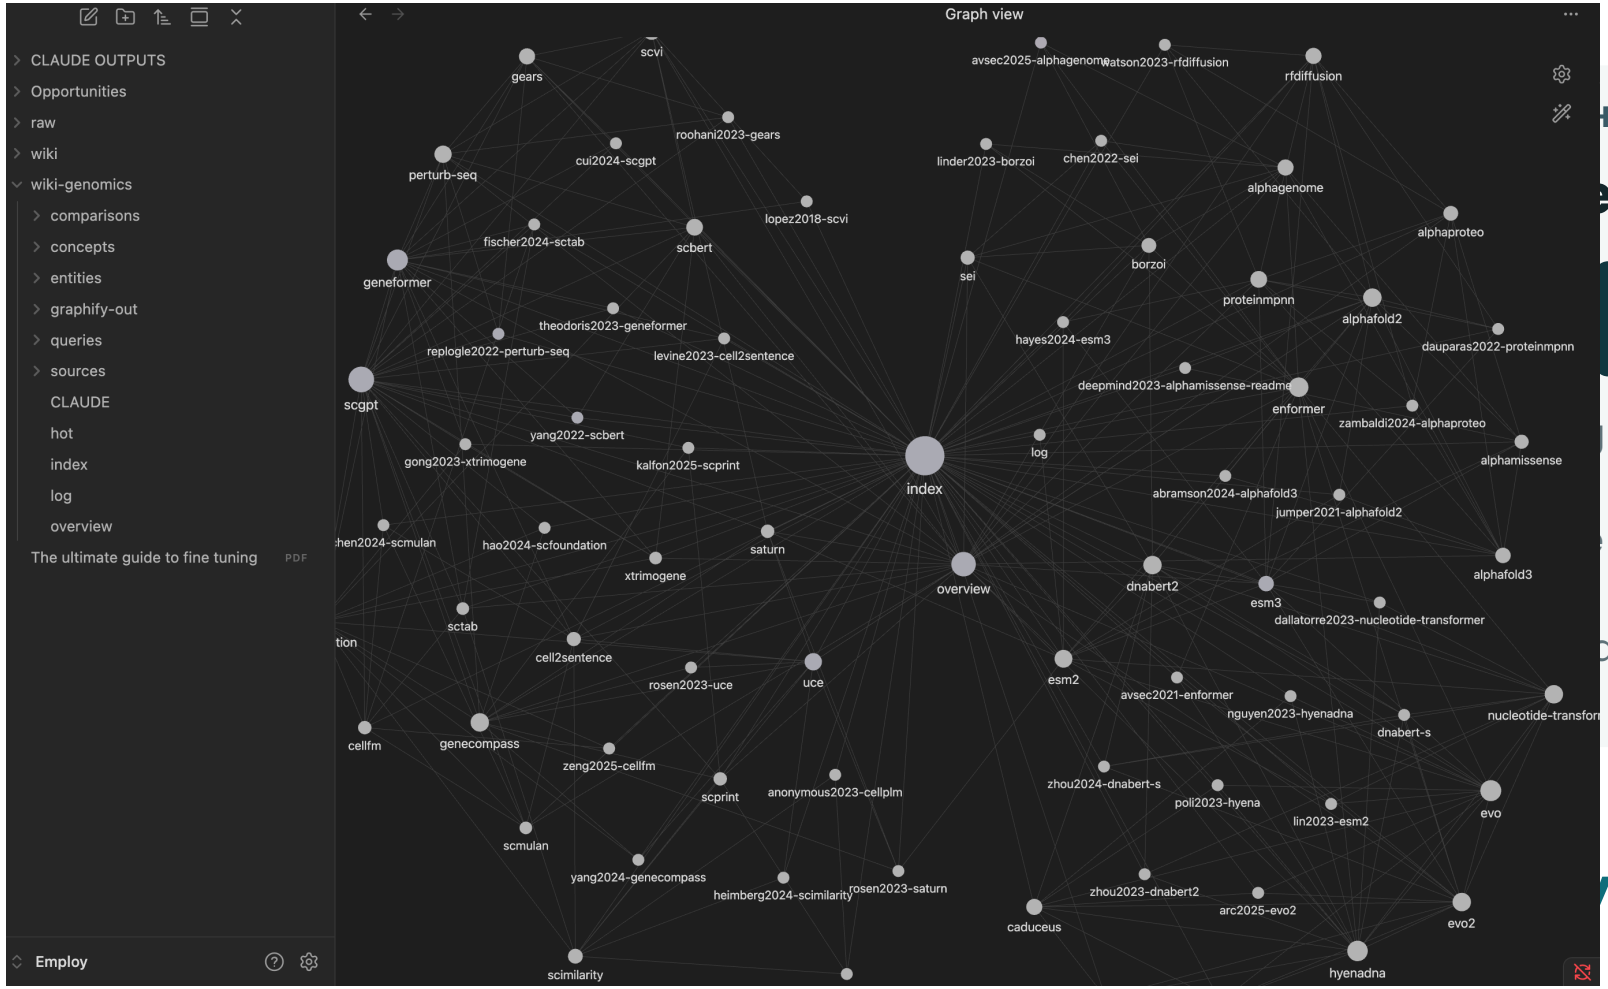
on a few GPUs,
, keep sensitive
control.

Sovereignty is

ach.

NVIDIA Blackwell (2024); Kigali Global AI Summit & Cassava Technologies GPU rollout (2025).

Knowledge management: Making AI know and NARO's decades-long KDB



H THE MODEL A SKILL

en-weight model

NARO labelled data → **Tuned model**

on local crop images or local-

and domain; runs offline once

ata and some GPU, but the

y answer cites a source.

RAG: Lewis et al. (2020). Knowledge-graph retrieval / GraphRAG: Edge et al. (2024).

THE FUTURE OF KNOWLEDGE

Human capital × token capital



Satya Nadella, Microsoft CEO

The institution of the future is a learning loop: people's judgment and the AI it owns, compounding together.

HUMAN CAPITAL

Judgment, know-how, relationships, design, intuition, and pattern recognition: **scientists & extension officers.**

captures judgment →

compounds ⟳

← *amplifies people*

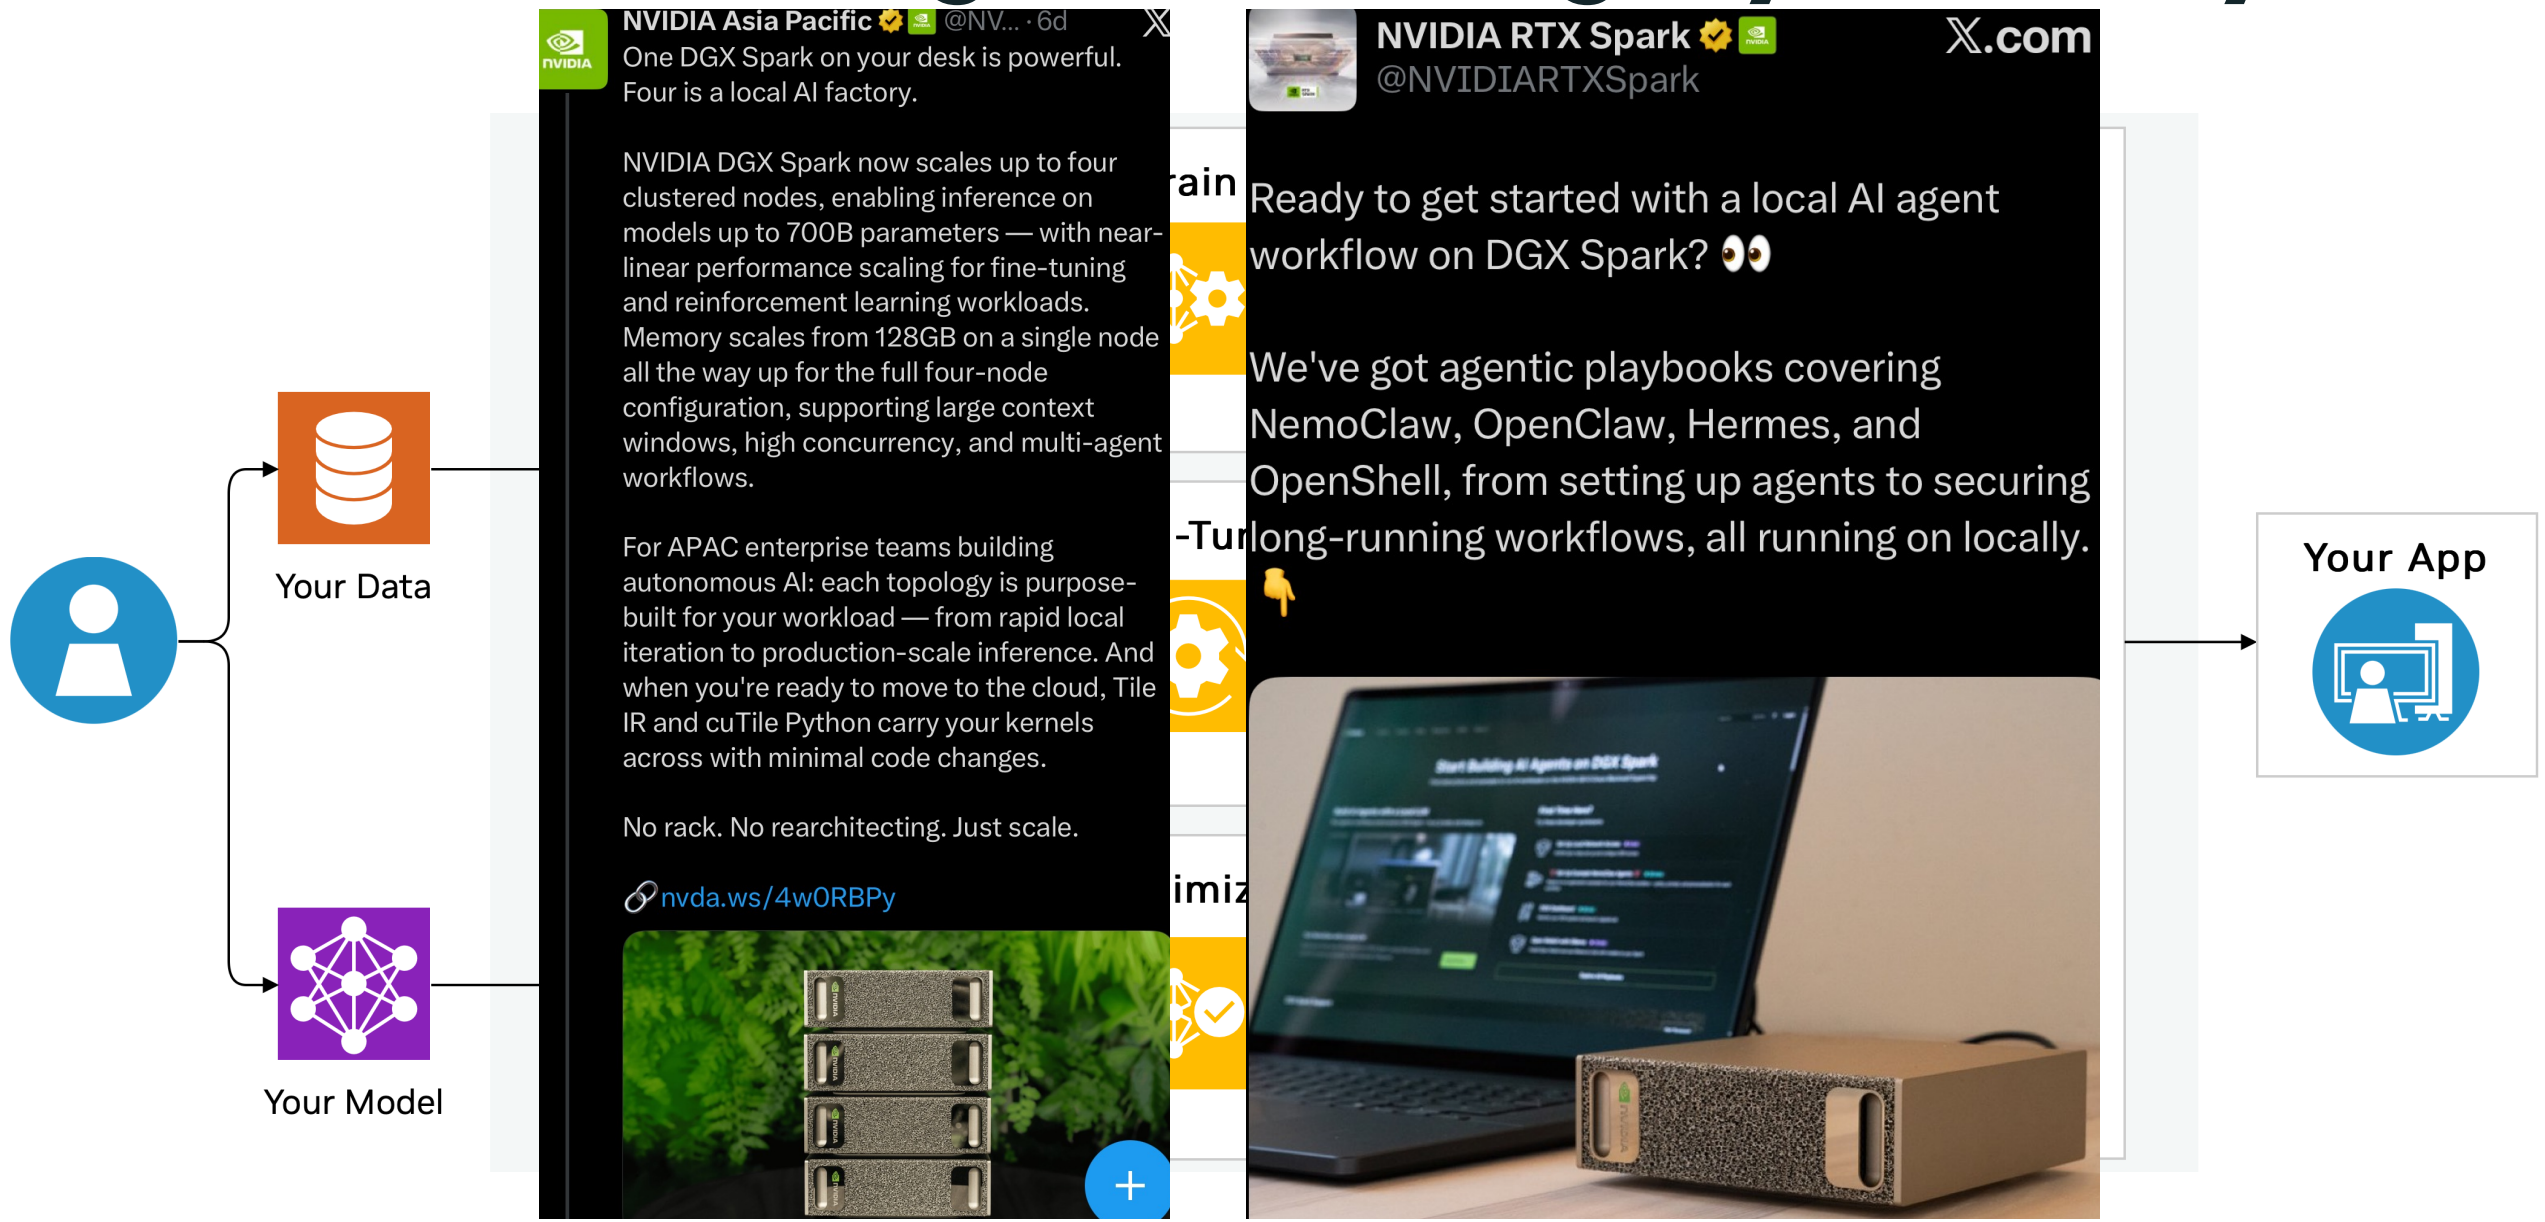
TOKEN CAPITAL

The AI capability you build and own: your models, your RAG archive, your fine-tuned skills.

Your edge is your people's knowledge, captured into systems you own rather than ceded to models you rent.

Framework after S. Nadella, "human capital & token capital" / the learning-loop firm (2026).

NVIDIA: making AI sovereignty a reality



A staged path to adoption..

1 AI Strategy

Set the goals.

Name the 2–3 research problems AI must solve, not the tools.

2 AI Plan

Pick the use cases.

Cassava genome surveillance, coffee breeding, the KM archive.

3 AI Ready

Prepare data and teams.

Digitise records; upskill early-career scientists.

4 Govern AI

Ensure responsible use.

Consent, provenance, access tiers, agreed up front.

5 Secure AI

Protect models and data.

Self-hosted where sensitive; sovereignty by design.

6 Manage AI

Keep systems reliable.

Monitor drift, version models, refresh the index.

AI: myths vs facts

	MYTH		FACT
✗	AI is a brand-new technology	✓	Coined in the 1950s; its foundations are decades old.
✗	AI will replace entire jobs	✓	It automates tasks and reshapes work; new roles emerge too.
✗	Computers can't be biased	✓	AI inherits and can amplify bias in its training data.
✗	AI proves cause and effect	✓	It finds correlation, not causation. Verify before acting.
✗	AI needs no human oversight	✓	Human oversight stays essential at every stage.

Adopt it with realistic expectations: AI is powerful, but it still needs human in the loop control or tight system rules.

Bipartisan Policy Center, "AI: Facts and Myths" (2026).

AI will not replace the scientists, only those who use it will out-pace those who don't.

Acknowledgments

Master students
BSc. Tahmina Begum
BSc. Lisa Glantschnig
Bsc. Leonard Jank



