Social X Factor

Handwerk der Zukunft

Maximilian Dauner
maximilian.dauner0@hm.edu
Munich University of Applied Sciences
Munich Center for Digital Sciences and AI (MUC.DAI)

AGENDA

HM MUC.DAI
Munich Center for
Digital Sciences and AI

The Craft of Entrepreneurship Is Changing

How AI is reshaping the way we discover opportunities, develop ideas, and build new ventures.

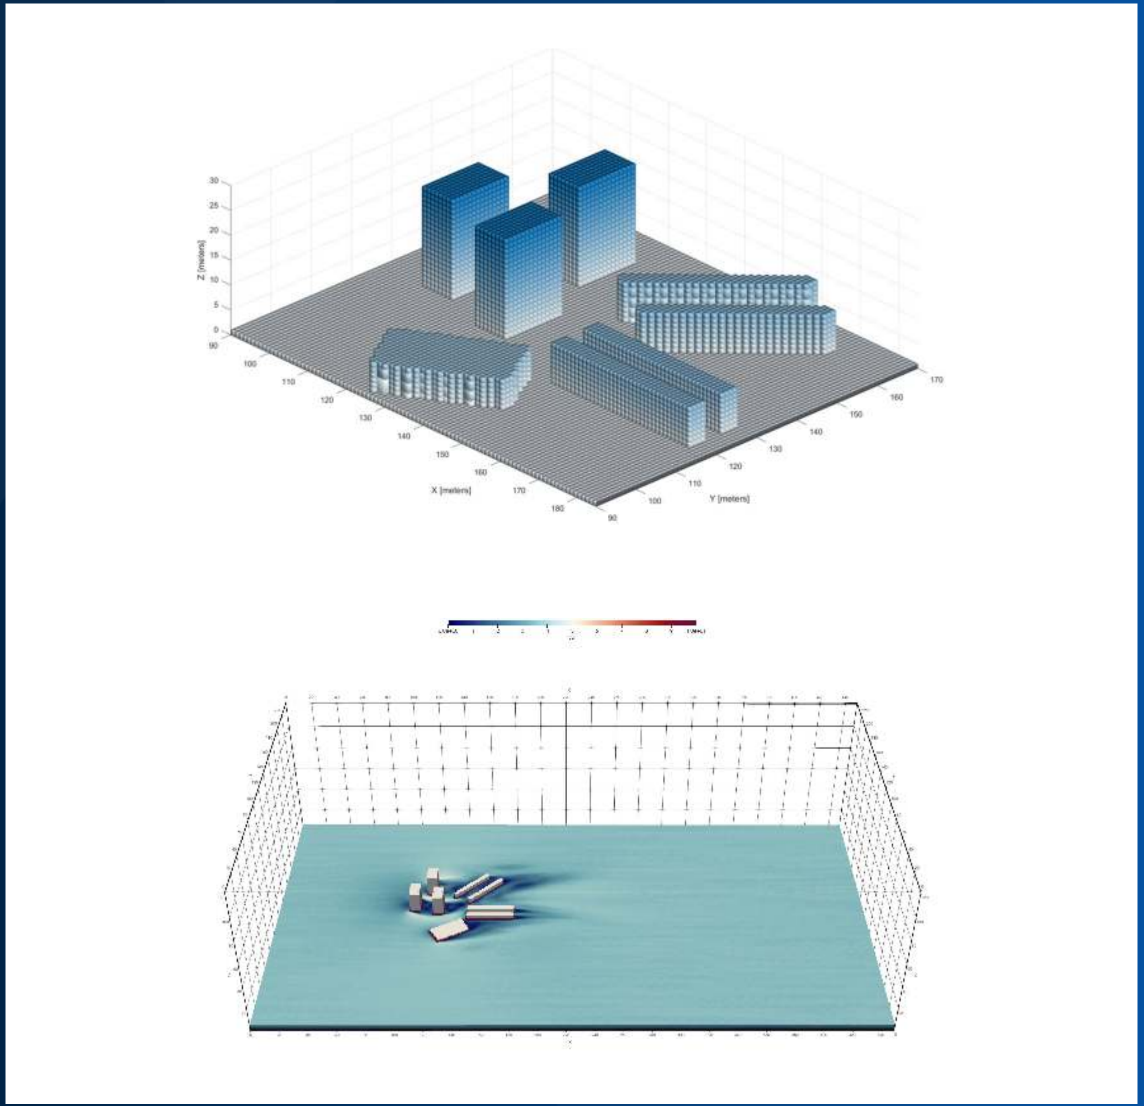

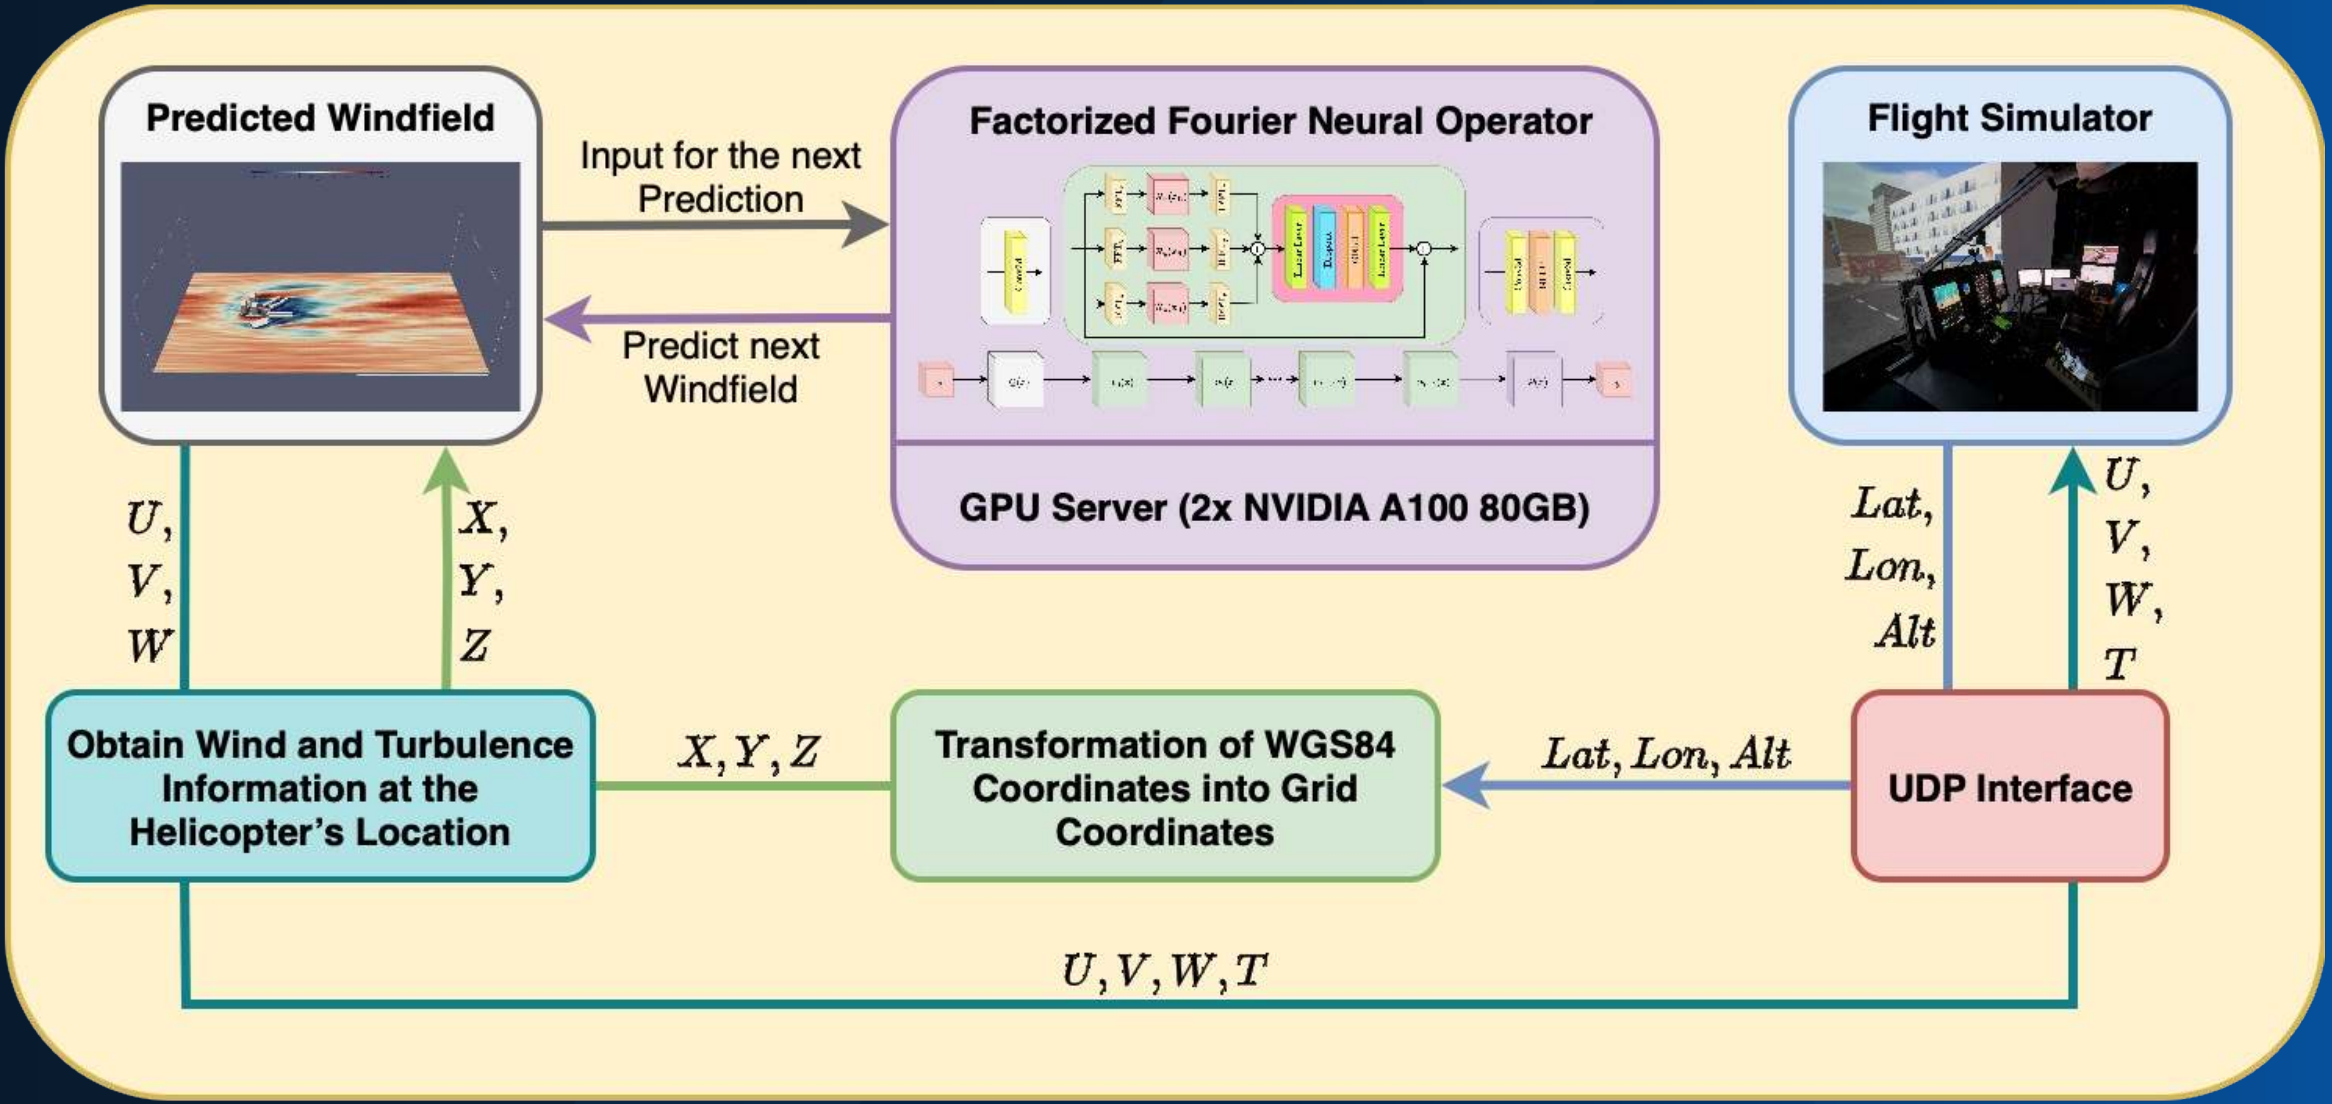
Predicted Windfield
Input for the next Prediction
Predict next Windfield
Factorized Fourier Neural Operator
GPU Server (2x NVIDIA A100 80GB)
Flight Simulator
U, V, W
X, Y, Z
Lat, Lon, Alt
U, V, W, T
Obtain Wind and Turbulence Information at the Helicopter's Location
X, Y, Z
Transformation of WGS84 Coordinates into Grid Coordinates
Lat, Lon, Alt
UDP Interface
U, V, W, T

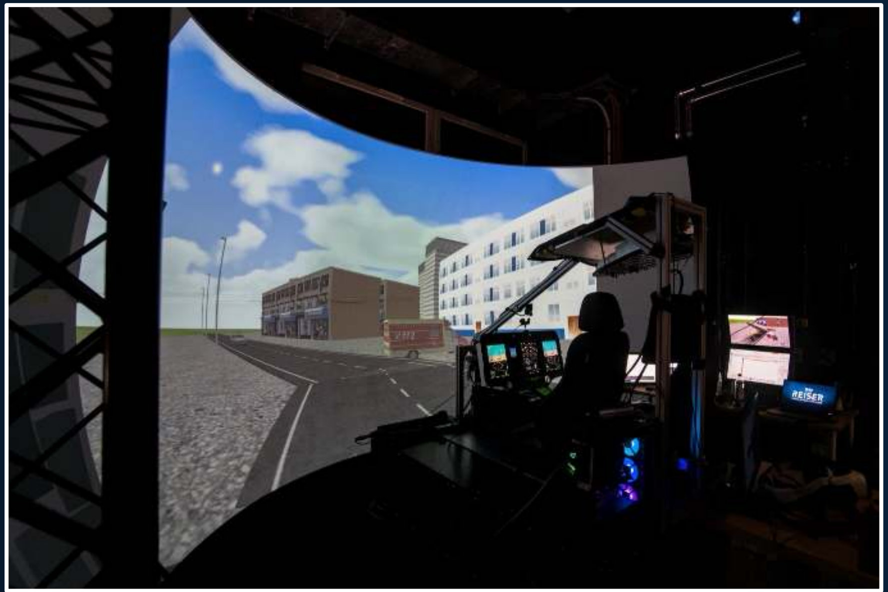


REISER

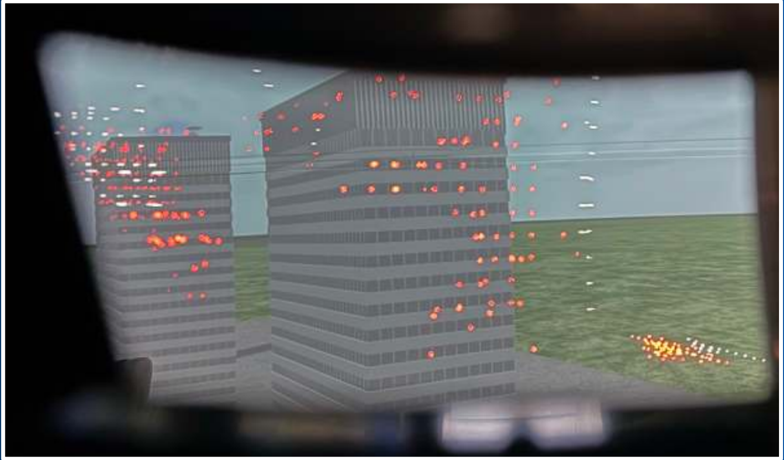

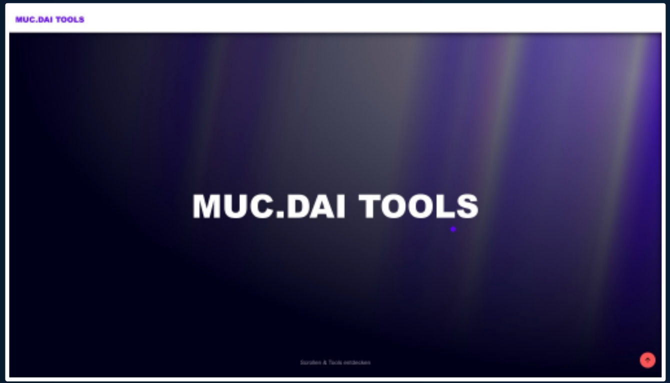

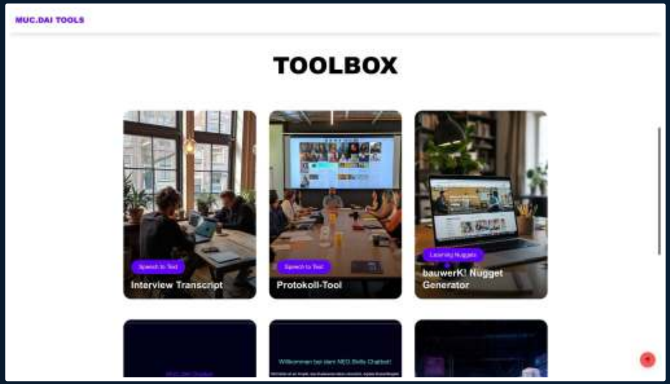

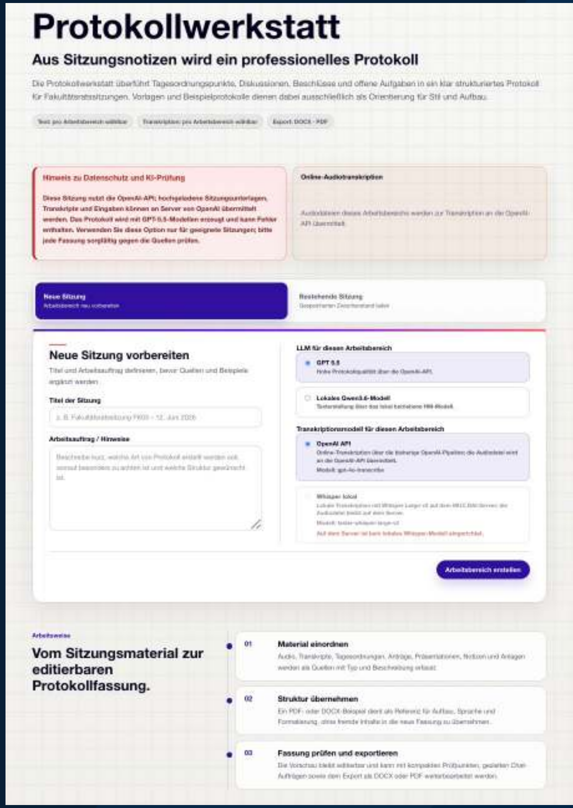


Beyond Note-Taking: Empowering Examiners with a Diarization-Aware and Transparent AI Assistant for Oral Exams

Maximilian Dauner, *MUC.DAI*, *University of Applied Sciences Munich*, Munich, Germany, maximilian.dauner0@hm.edu

Ricky Liu, *MUC.DAI*, *University of Applied Sciences Munich*, Munich, Germany, ricky.liu@hm.edu

Gudrun Socher, *MUC.DAI*, *University of Applied Sciences Munich*, Munich, Germany, gudrun.socher@hm.edu

Faster Than Light × XPENG

A Creative AI Production System for the Next Generation of Mobility.

Prepared by Maximilian Dauner & Matteo Walter

July 2026 · Concept Pitch

Built on Cinema.

Our approach begins with the language of film: storytelling, composition, atmosphere, performance and the discipline of professional production.

FTL × XPENG — Concept Pitch · Maximilian Dauner & Matteo Walter — 3

Cinematic Thinking, rebuilt for AI.

For the Runway Big Ad Contest, we translated creative direction, visual storytelling and commercial production into an entirely AI-native film.

FTL × XPENG — Concept Pitch · Maximilian Dauner & Matteo Walter — 4

From Campaign to Capability.

We design the creative, technical and educational system that allows XPENG to create, learn and scale AI-driven content internally.

1. Discover
Define brand goals, content priorities and use-cases.

2. Foundation
Structured knowledge and asset layer.

3. Architect
Evaluate setups and define requirements.

4. Create
Build workflows for ideation, generation and campaigns.

5. Enable
Design internal interfaces for independent ideation and creation.

6. Transfer
Workshops, learning nuggets and tailored guidance.

FTL × XPENG — Concept Pitch · Maximilian Dauner & Matteo Walter — 7

Introduction to Artificial Intelligence

A first look at the concepts behind AI systems.



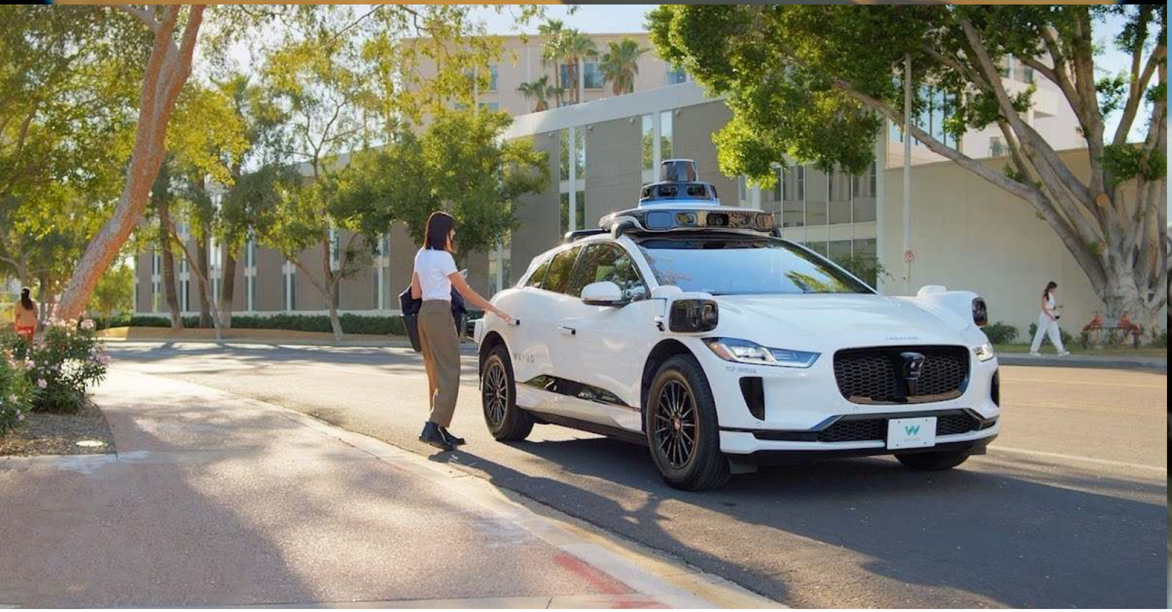

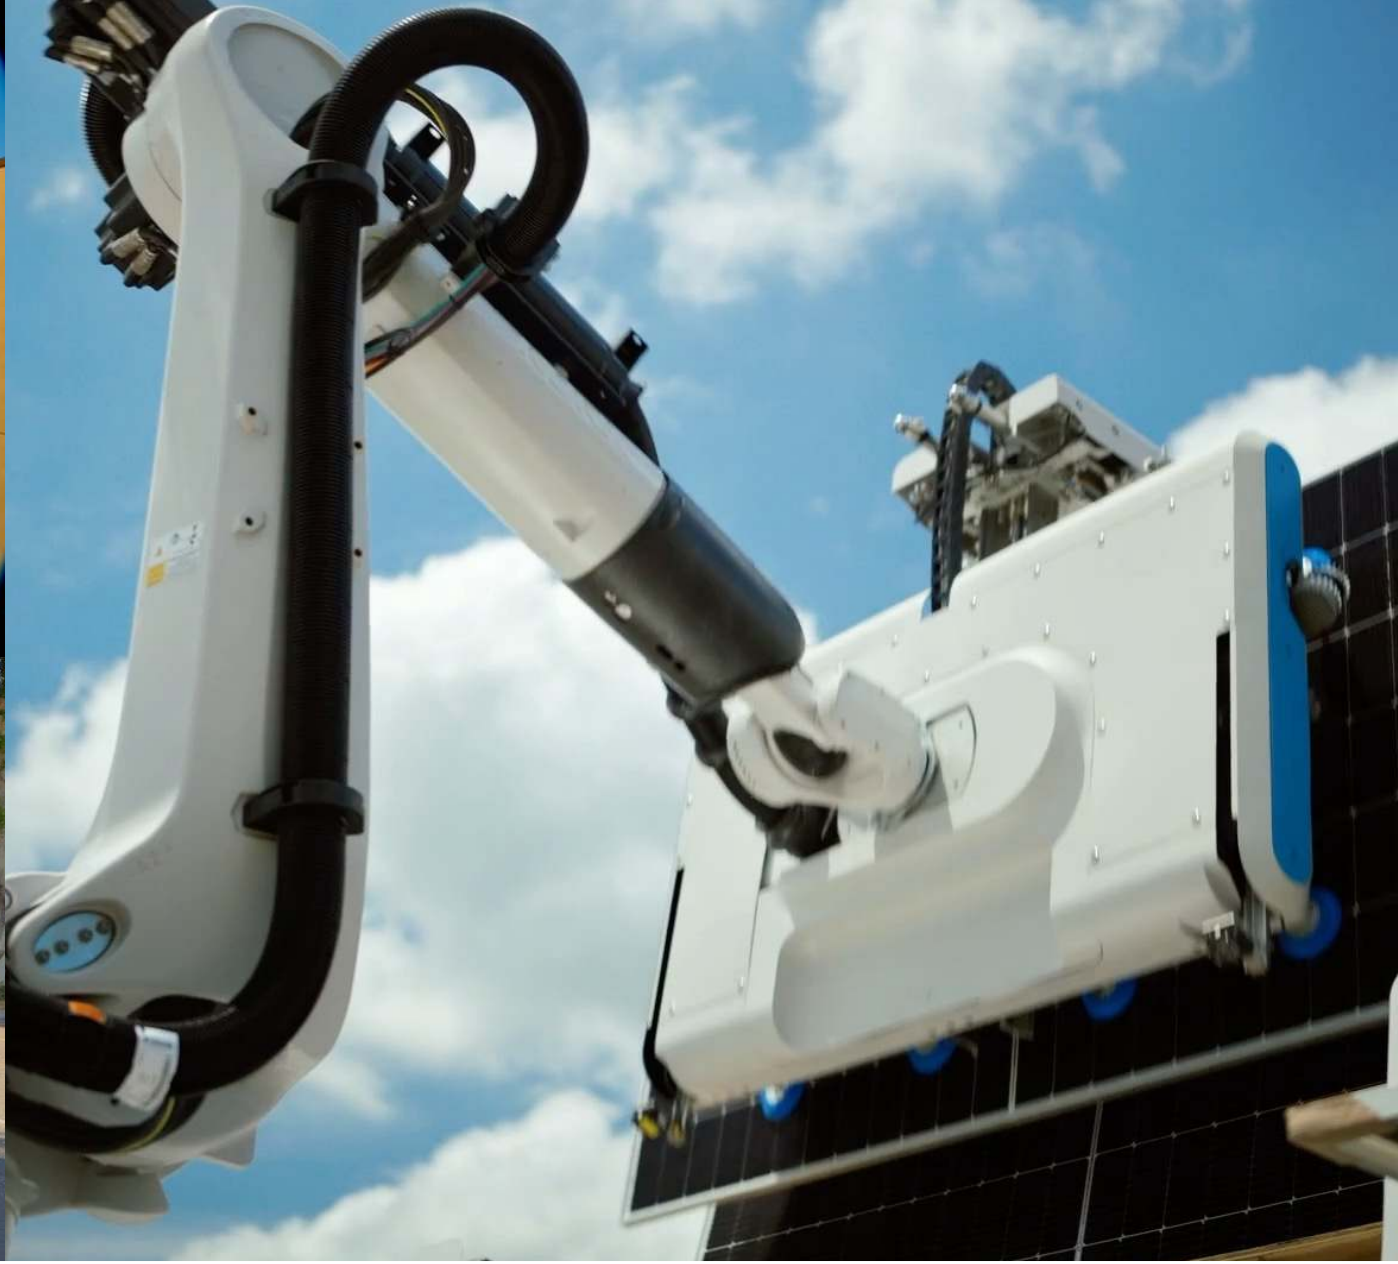

NEO

Home Robot

Product FAQ

Standard **$499/mo**
Subscription

- Monthly Subscription
- Starter Productivity Package
- Standard Delivery

Early Access **$20,000**
Ownership

- Ownership with 3-Year Warranty
- Premium Support
- Priority Delivery

$200 Deposit Due Today

Fully Refundable

US Deliveries start 2026

Order Now

Artificial Intelligence (AI)

Machine-based systems that infer from inputs how to generate outputs such as predictions, content, recommendations, or decisions.

Machine Learning (ML)

A subset of AI in which models learn patterns from data instead of relying only on manually defined rules.

Neural Network (NN)

A parameterized computational model composed of interconnected layers that learn transformations from input data.

Deep Learning

A subset of machine learning that uses neural networks with many computational layers to learn complex representations from data.

Gemini

Claude

Discriminative AI

Models that learn to predict a target output from an input, such as a category, score, probability, or numerical value.

They make decisions by identifying differences between classes and determining which category an input most likely belongs to.

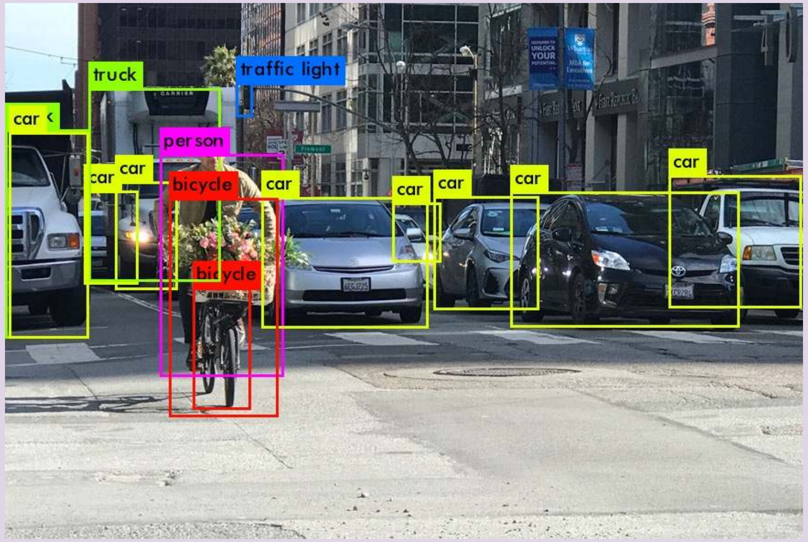

Generative AI

Models that learn statistical patterns from data and use them to generate new text, images, audio, video, or code in response to an input.

They learn patterns from existing data and use them to generate new, similar, but original examples.

Dataset

A collection of examples used to train, adapt, validate, or evaluate a model.

Model

A computational system with learned parameters that maps inputs to outputs.

Training

The process of adjusting a model's parameters using data and an objective function.

Inference

The process of using a trained model to generate a prediction, decision, or other output for an input.

Dataset

A collection of examples used to train, adapt, validate, or evaluate a model.

Model

Goal: Human Face

Dataset

Model

Goal: Human Face

A collection of examples used to train, adapt, validate, or evaluate a model.

Dataset

Model

Goal: Human Face

A collection of examples used to train, adapt, validate, or evaluate a model.

A Green Monkey

Model

Dataset

Goal: Text to Image

A collection of examples used to train, adapt, validate, or evaluate a model.

A Green Monkey

Dataset

A collection of examples used to train, adapt, validate, or evaluate a model.

Goal: Text to Image

Model

A snowy suburban street with a red car parked on the side and a green truck driving down the road. The street is lined with houses and trees, and there is a sign on the left side of the image

A Green Monkey

Dataset

A collection of examples used to train, adapt, validate, or evaluate a model.

Goal: Text to Image

Model

A red helmet with a black chin strap and a black and red logo on the side is displayed against a white background.

A Green Monkey

Dataset

A collection of examples used to train, adapt, validate, or evaluate a model.

Goal: Text to Image

Model

A man in a suit is standing at a podium with a microphone in front of him. He is wearing a light blue tie with a subtle pattern. The background shows a blurred view of trees and a clear sky.

Dataset

jpg dict	json dict
	{"prompt": "A promotional image for a crampon storage bag with a group of hikers in the background. The bag is prominently displayed in the foreground, with a large, detailed image of the bag's mesh design an…
	{"prompt": "A small, shiny copper-colored object is placed on a white surface."}
	{"prompt": "A stack of pancakes with a glossy, caramel-like syrup on top is presented. The pancakes are golden brown and appear to be freshly cooked. The syrup is dripping down the sides of the pancakes,…
	{"prompt": "The first image shows a close-up of a motorcycle's rear fender with the Harley-Davidson logo prominently displayed. The logo is centered on the fender, which is a dark color with a glossy finish.…
	{"prompt": "Aerial view of a tropical island with a long, curved row of traditional thatched-roof huts on stilts over the water. The huts are connected by a wooden walkway. The island has a sandy beach with a…
	{"prompt": "Aerial view of a lush green mountainous landscape with a winding path leading through the hills. The terrain is covered with dense vegetation, and the sky is partly cloudy."}

Dataset

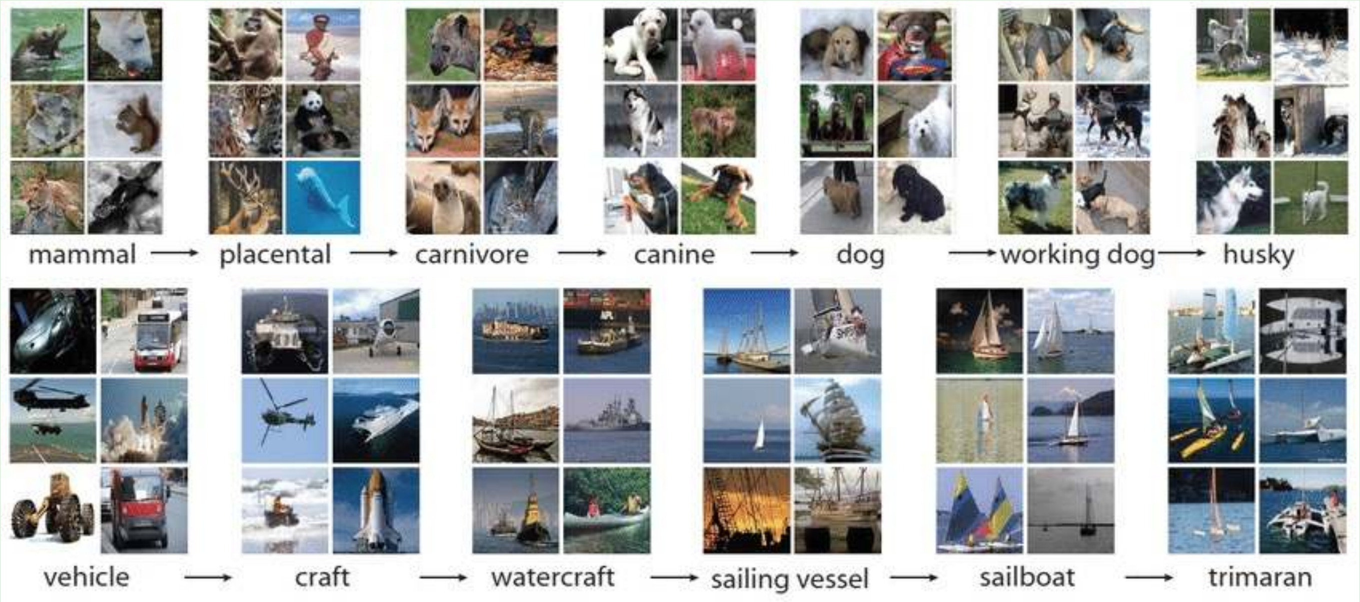
mammal → placental → carnivore → canine → dog → working dog → husky
vehicle → craft → watercraft → sailing vessel → sailboat → trimaran

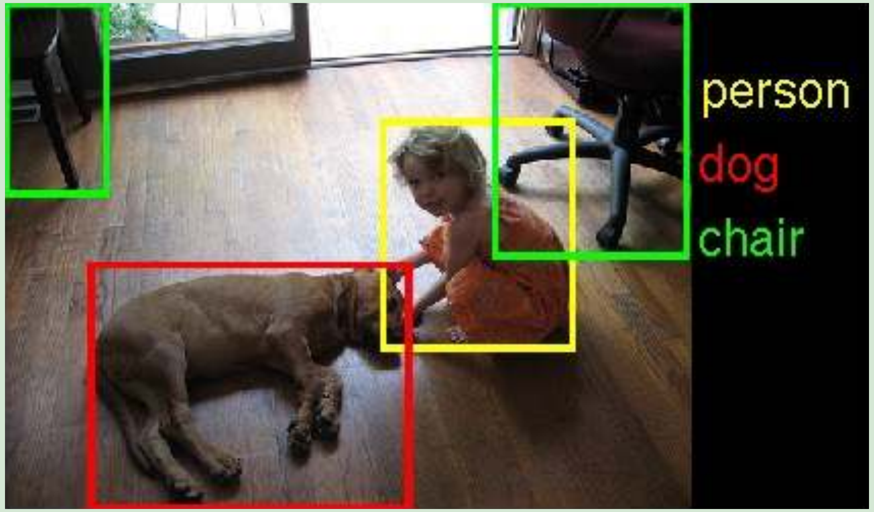
person
dog
chair

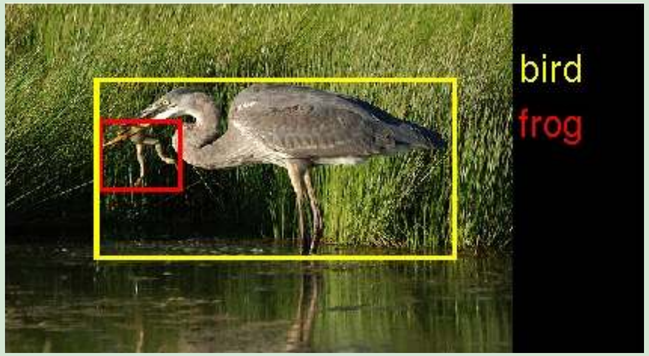
bird
frog

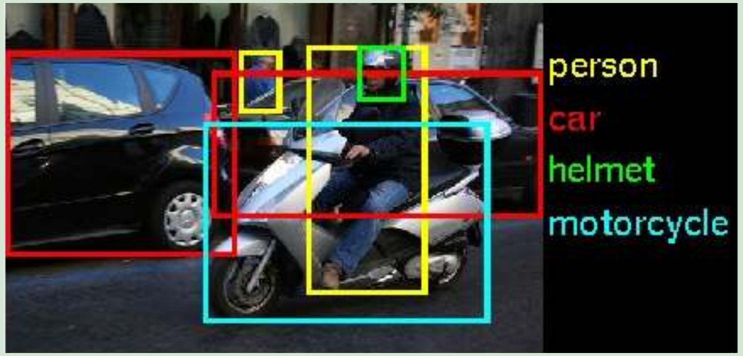
person
car
helmet
motorcycle

What does the model never see?

A model can only learn from the examples, patterns, and perspectives contained in its dataset.




NOMOS

TACHYMETER
Breitling

STERNGLAS
ZEITMESSER
AUTOMATIK


CONTRACT

Dataset

A collection of examples used to train, adapt, validate, or evaluate a model.

Model

A computational system with learned parameters that maps inputs to outputs.

Training

The process of adjusting a model's parameters using data and an objective function.

Inference

The process of using a trained model to generate a prediction, decision, or other output for an input.

The Hidden Costs of AI

AI may feel digital, but training and inference depend on physical infrastructure, electricity, water, and hardware.

Number of Parameters

The number of parameters define how much neural network can learn.

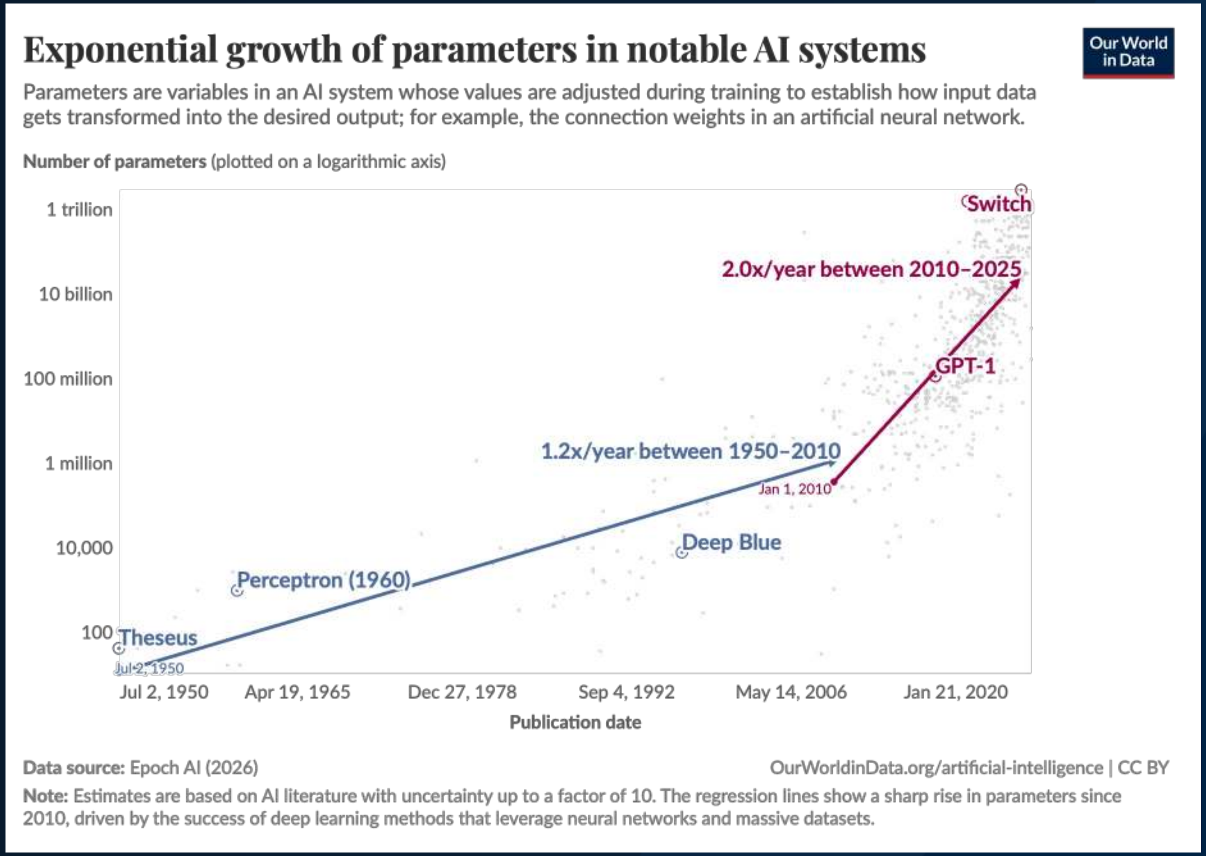

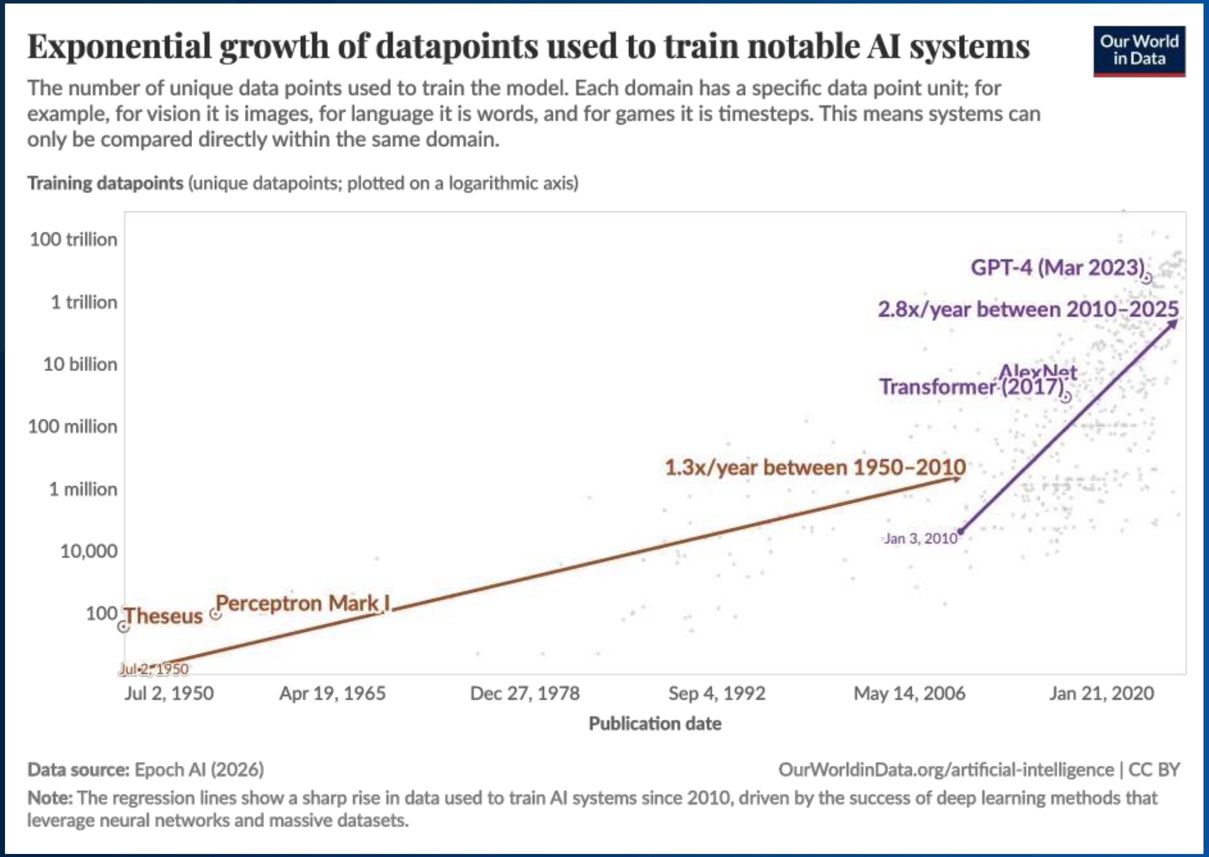

More Parameters, More Emissions

Larger models require more computation, resulting in greater energy use and carbon emissions.

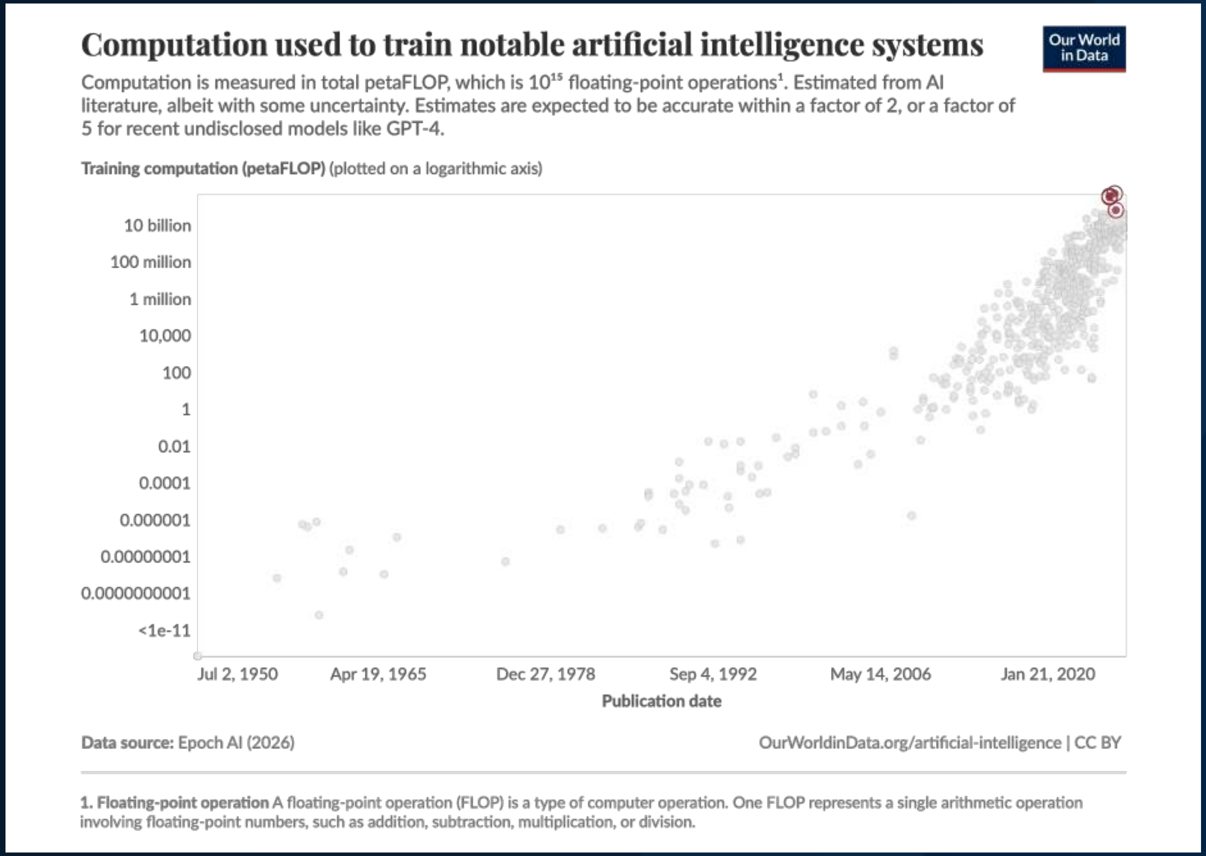

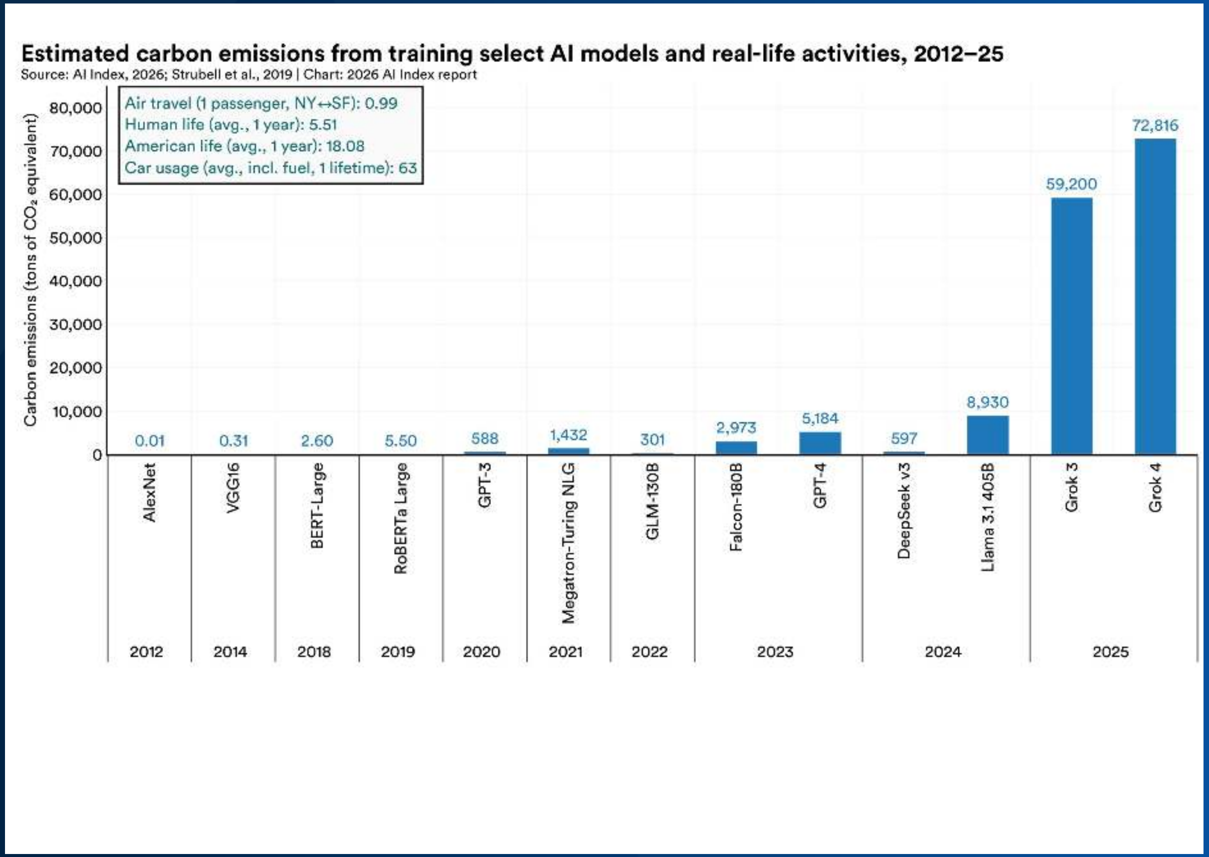

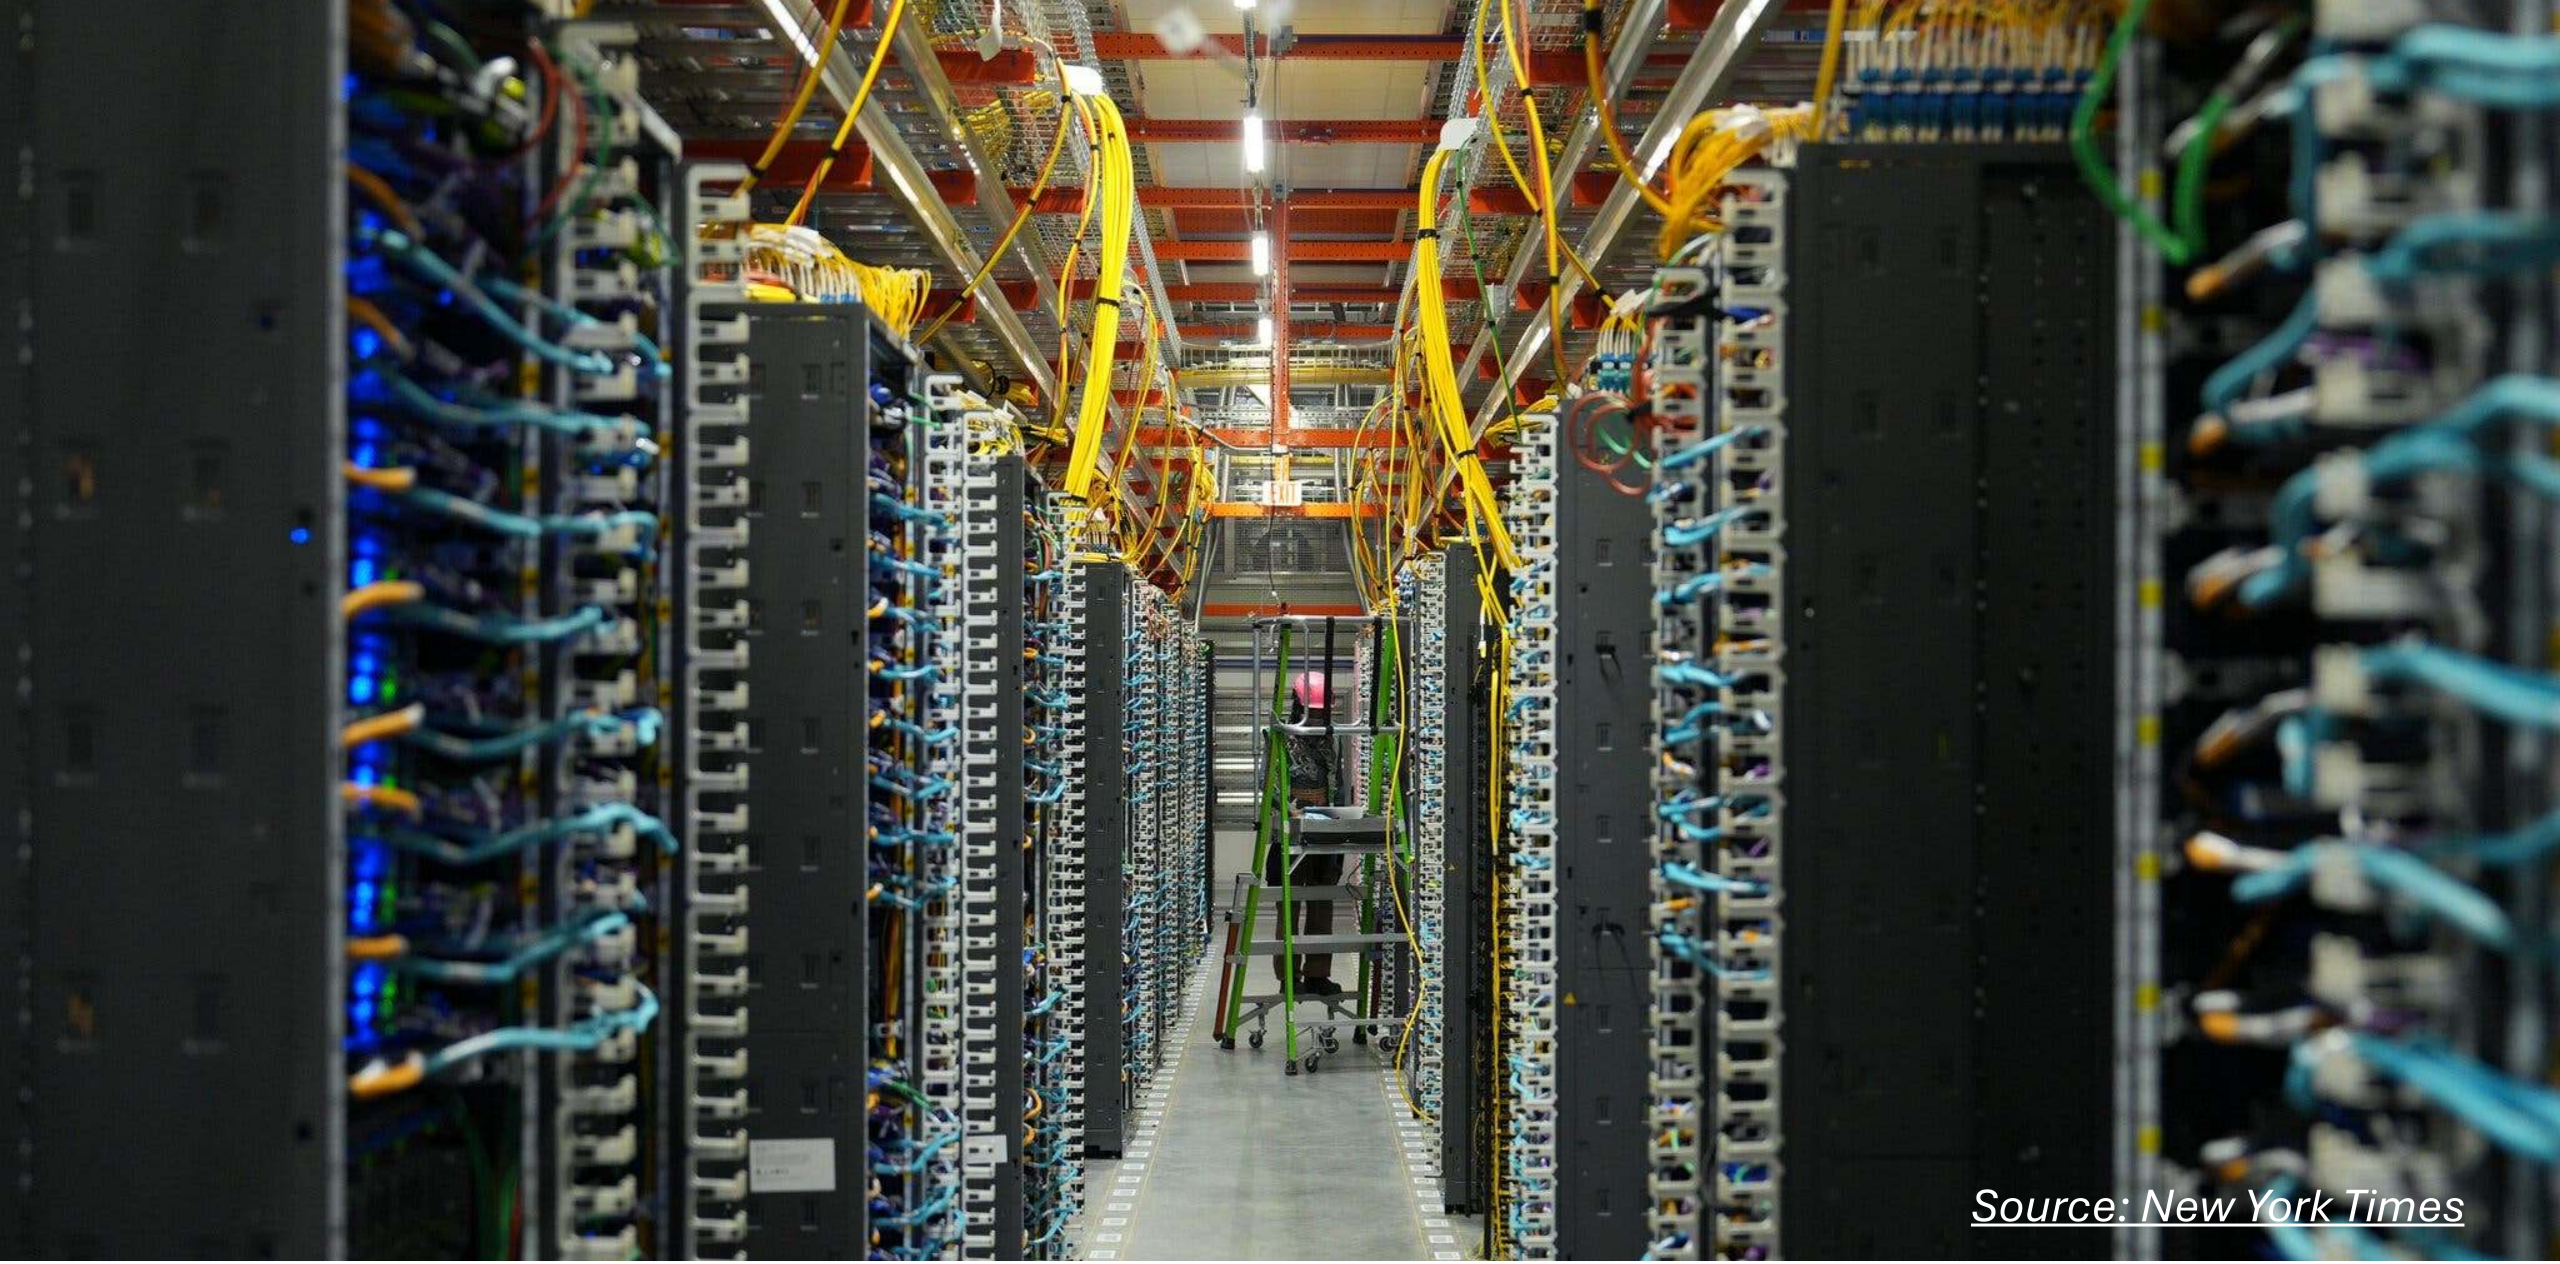

Source: New York Times

Source: New York Times

Global data center electricity demand could more than double by 2030

Rapid growth in AI and digital services is placing increasing pressure on global energy systems.

485 → 950 TWh
Global data-centre electricity consumption2025 → 2030 (roughly all of Japan's consumption)

Around 3%
Share of global electricity demand projected for 2030

3×
Projected growth in electricity use by AI-focused data centres, 2025–2030

Source: IEA "Energy and AI" / The Guardian

Source: New York Times

Source: New York Times

$130 billion of data center projects were blocked or delayed by local opposition — in Q1 2026 alone. Full-year 2025 already reached $156 billion, up from an initial $64 billion estimate.

Data Center Watch, Q1 2026 report

From Impact to Responsibility

The question is not only what AI can do, but what it should do—and who remains accountable.

What Do We Mean by "AI Ethics"?

Exploring how values, rights and responsibilities shape the design and use of AI systems.

Fairness

Does the system treat people equitably across race, income, ability, and language?

Transparency

Can users and regulators understand how the system reaches its outputs?

Privacy

Is personal data collected, used, and retained only as necessary?

Accountability

When something goes wrong, who is responsible — and can it be traced back to them?

Human oversight

Can a person meaningfully review, override, or halt the system?

From Understanding AI to Building With It

Now we translate technical possibilities, environmental limits, and ethical responsibilities into entrepreneurial action.

An AI-Native Workflow Connects More Than a Prompt

Reliable outcomes emerge from context, tools, memory, evaluation, and human review.

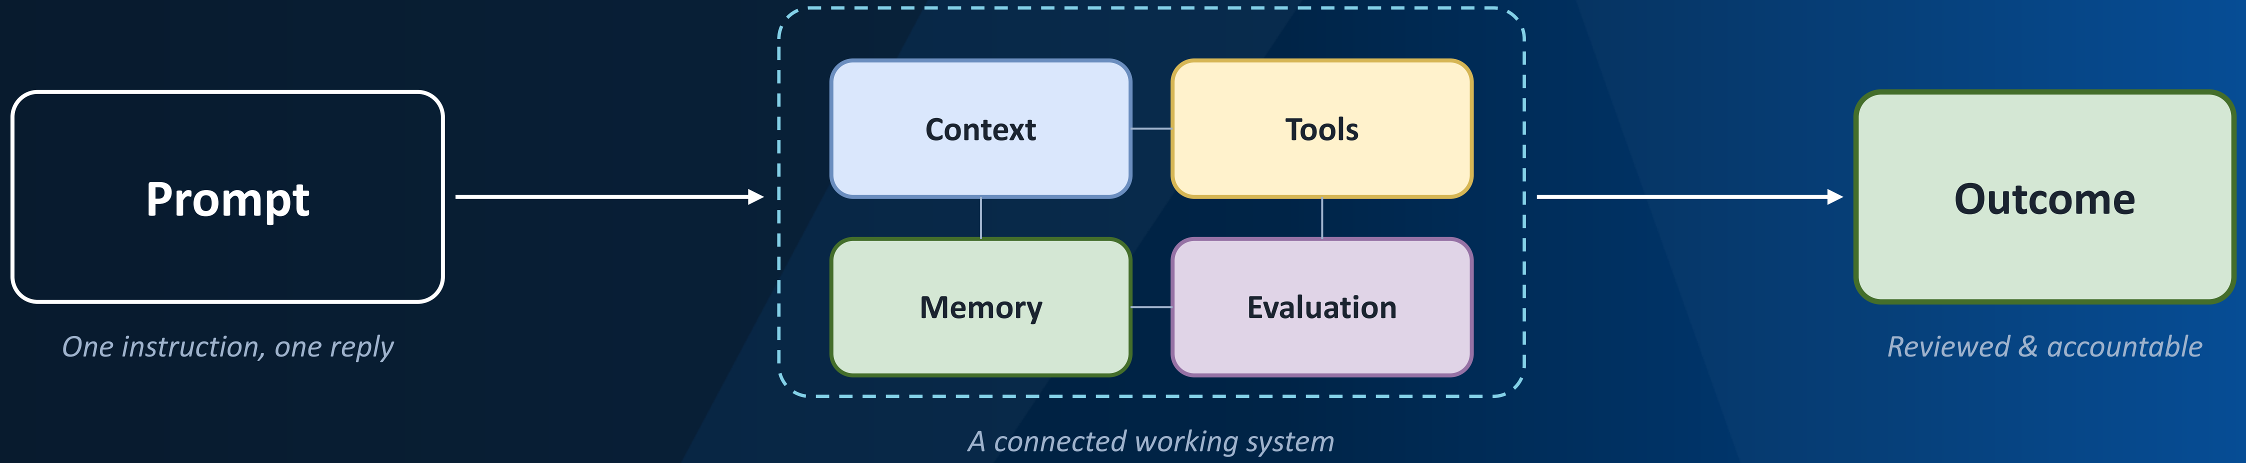

A Good Prompt Is a Task Brief, Not a Spell

Clear goals, context, constraints, output requirements, and success criteria improve an interaction—but they do not create a complete process.

PROMPT

Outcome	**Context**	**Task**
Constraints	**Output**	**Success Criteria**

Necessary for good instructions. Insufficient for complex work.

From Prompt Engineering to Context Engineering

Reliable AI work depends on what the system knows, can access, remembers, and is allowed to do.

Vision and Intent	What are we trying to achieve?
Context and Evidence	What should the system know and trust?
Workflow and Tools	Which steps and capabilities are required?
Memory and State	What must persist across sessions?
Evaluation and Governance	How do we know it worked, and who decides?

The Entrepreneur Becomes a Workflow Designer

Humans define the vision, evidence, boundaries, and decisions; AI systems help explore, execute, compare, and iterate.

TRADITIONAL USE

- Human researches
- Human drafts
- Human analyzes
- Human builds
- AI improves one artifact

AI-NATIVE WORKFLOW

- Human defines the goal
- AI gathers & structures information
- Tools execute research, analysis, or coding
- Evaluators test the result
- Human interprets and decides

Direction | Evidence | Judgment | Accountability

Build a Shared Source of Truth

Long chats are temporary; project knowledge should live in reusable, versioned, and searchable artifacts.

```
venture-workspace/
  • 00_north_star/
    • vision.md
    • decision_to_make.md
  • 01_customers/
    • segments.md
    • interview_guide.md
    • transcripts/
  • 02_evidence/
    • research_log.md
    • evidence_matrix.csv
  • 03_assumptions/
    • assumption_backlog.md
  • 04_product/
    • requirements.md
    • acceptance_tests.md
    • style_guide.md
  • 05_decisions/
    • decision_log.md
  • 06_ai_system/
    • agent_instructions.md
    • evaluation_cases.csv
```

One canonical project state—accessible to humans and AI.

Markdown-based vaults (e.g. Obsidian) are one option among several.

Capture Thinking Before It Disappears

Record discussions, transcribe them, and convert them into assumptions, decisions, open questions, and next actions.

Discussion / voice note → Consent-based recording → Transcription → **Extraction agent** → Human review → Canonical artifacts

Extraction agent ↓

Assumptions | **Decisions** | **Open Questions** | **Actions & Owners**

Observation ≠ Interpretation ≠ Assumption ≠ Decision

Choose the Simplest System That Works

Use chat for one-off tasks, workflows for repeatable steps, and agents only when the path must adapt.

Multi-Agent System
Independent specialization creates measurable value
Example: researcher + critic + synthesizer

Single Agent
The system chooses tools and next steps
Example: research across varied sources

Workflow
The sequence of steps is known in advance
Example: transcribe → summarize → classify

Chat
One response is sufficient
Example: rewrite a value proposition

Do not build a multi-agent system because it looks advanced.

Deep Research Is a Process, Not a Prompt

Start with a decision and a source policy, then plan, search, compare, challenge, synthesize, and verify.

Decision to Support → Research Plan → Source & Evidence Policy → Parallel Investigation → Evidence Matrix → Contradiction & Gap Analysis → Synthesis → **Human Decision**

WEAK REQUEST

"Research the food-waste market."

STRONG BRIEF

Decide: enter Munich food-waste market? · Geography: Munich metro
Sources: industry reports, interviews, gov data · Timeframe: last 3 years

Claim	Evidence	Source	Date	Confidence	Contradiction
Munich restaurants discard food weekly	Interview + industry report	Industry association	2025	Medium	None found
Donation is blocked by food-safety rules	Regulation review	EU / German rules	2024	High	Contradicts claim above

Agentic Market Discovery

Let agents gather and structure evidence across sources—then let humans decide which opportunity is worth pursuing.

INPUT

- Venture vision
- Target group
- Geography
- Timeframe
- Current assumptions
- Approved sources
- Output format

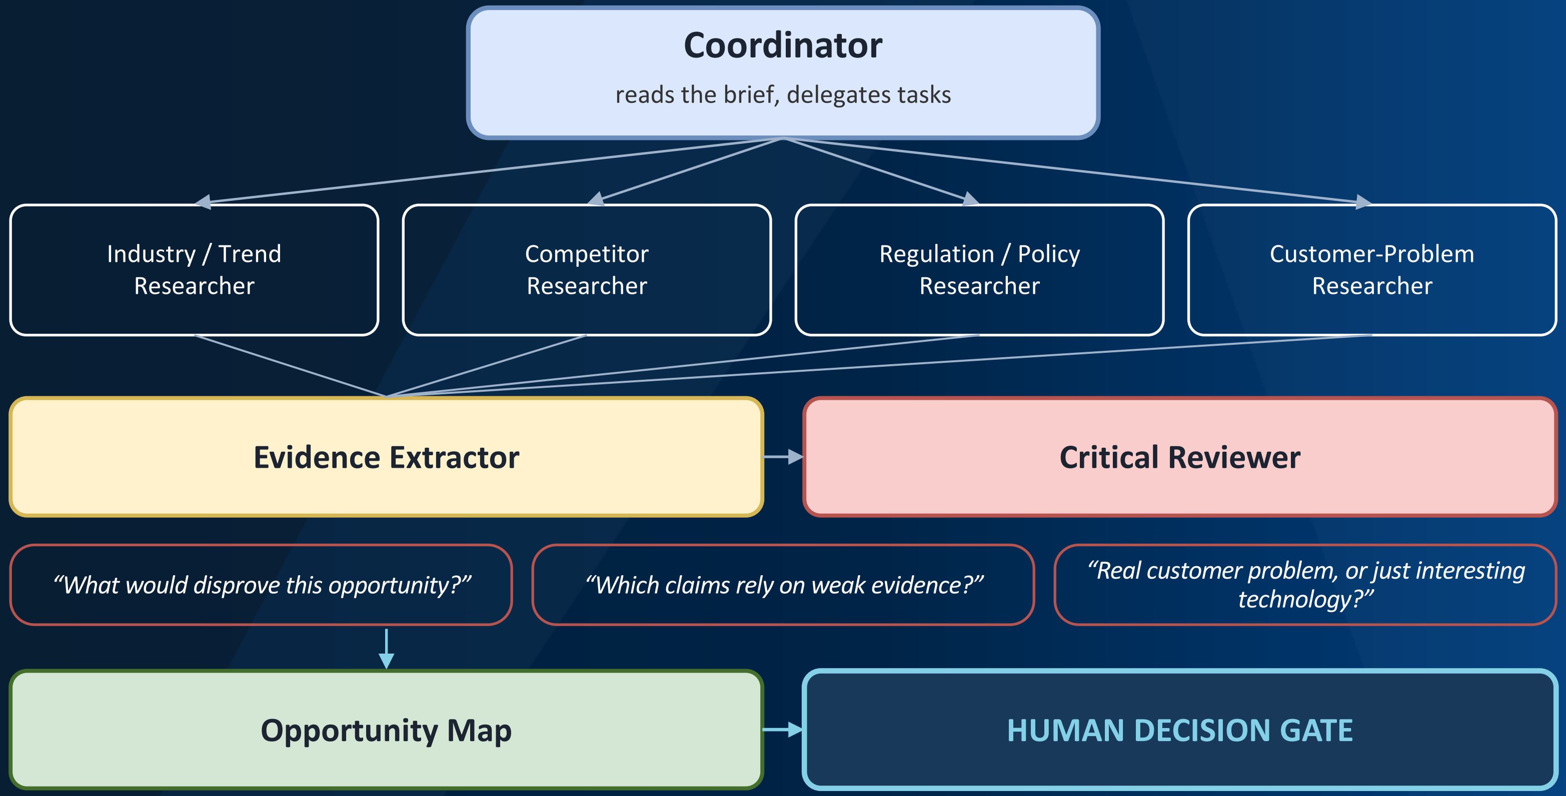

OUTPUT

- Evidence matrix
- Competitor map
- Opportunity areas
- Contradictions
- High-risk items
- Recommended experiments

Synthetic Customers are not an Evidence

Simulated users can test questions and reveal possibilities, but real customers remain the source of validation.

SIMULATION

Existing Research / Real Interviews

Grounded Customer Profiles

Interviewer Agent ⟷ **Simulated Respondent**

Gets: topic guide, anti-leading rules, adaptive follow-up

Gets: background, goals, constraints, uncertainties

Interview-Quality Critic

Possible Needs / Revised Questions

Output: Hypotheses, not findings.

VALIDATION

Revised Interview Guide

Real Participants

Recorded / Transcribed Interviews

Human-Validated Findings

Turn Interviews Into a Learning System

A repeatable pipeline transforms conversations into traceable themes, hypotheses, experiments, and product decisions.

Recorded Interview → Timestamped Transcript → Quote / Observation Extraction → Qualitative Coding → Theme Clustering → Contradiction / Outlier Analysis → Needs / Opportunity Statements → **Human Validation** → Backlog & Next Experiments

TRACEABILITY CARD

Finding	Supporting quotation	Participant	Timestamp	Interpretation	Confidence

Separate: explicitly said → observed → inferred → unknown

Vibe Coding Needs a Specification

Natural-language coding becomes reliable when the vision is translated into requirements, acceptance tests, constraints, and review cycles.

Vision → Specification & User Stories → Acceptance Criteria → Agent Explores & Plans → Implementation in Sandbox / Branch → Automated Tests & Visual Inspection → User Feedback → Updated Specification

WEAK INSTRUCTION

“Build an application for restaurant food waste.”

BETTER SPEC

User: independent restaurant manager, <5 min/day · Flow: select meal → enter quantity/window → publish (≤3 taps)

Acceptance: mobile-friendly · blocks missing quantities · offers auto-expire · critical flows tested · no personal data

“Agentic prototyping” = structured, spec-driven vibe coding.

Use Code to Create Verifiable Evidence

For analyses and plots, preserve the data, transformations, assumptions, code, and figure—not only the image.

Raw Data → Cleaning / Validation Script → Analysis Code → Reconciliation / Statistical Checks → Plot & Summary Table → Caption, Assumptions & Limitations

`raw_data.csv` | `clean_data.csv` | `analysis.py` | `results.csv` | `figure.svg` | `methods.md`

A beautiful plot is not evidence unless the path from data to figure is inspectable.

Agents Need Boundaries and Tests

Autonomy is useful only when actions are observable, reversible, evaluated, and approved at the right moments.

Action	Control
Read public information	Automatic
Create an internal draft	Automatic, with trace
Change a shared document	Review or version history
Contact a customer	**Human approval**
Publish externally	**Human approval**
Financial / legal / sensitive action	**Human decision**

Redesign One Entrepreneurial Workflow

Ask what should be captured, delegated, verified, automated, and kept human.

Goal and Decision

What are we trying to decide?

Current Process

How is this done today, step by step?

Context

What information and evidence are needed?

AI and Tools

Where could AI add real value?

Human Gates

Where must a person approve or decide?

Evaluation

How will we know it actually worked?

Do not begin with an AI tool. Begin with a problem, a decision, and a measurable outcome.

Thank you for your attention!

Social X Factor – Handwerk der Zukunft

Maximilian Dauner

maximilian.dauner0@hm.edu

Munich University of Applied Sciences

Munich Center for Digital Sciences and AI (MUC.DAI)