Can AI Agents Use Our Agricultural Research?

CAST Webinar

09/22/2026

Key takeaway

- Industry professionals and researchers increasingly use **AI agents** (AI systems such as ChatGPT or Claude that can search, run tools, and act on a request, not only chat) to find and apply scientific findings to agriculture related issues and questions
    - Need to rethink how we deliver research outcomes.
- Today's principles for sharing data (**FAIR**: Findable, Accessible, Interoperable, Reusable) and research software (**FAIR4RS**) were written for human and machine. Making research **agent-actionable** (usable by an AI agent without) needs additional standards for discovery, execution, interpretation, and traceability.
- We show, with one scientific modeling paper, how a paper can become a **knowledge unit**: a package of model, data, and instructions that anyone with an AI agent can reuse.

Outline

- Background
  - Impact of AI on science
- Challenges
  - Scientific knowledge is scattered
  - Scientific models can be hard to reuse
- FAIR4AG2 principles
- A case study
  - A plugin for milk spoilage decision-support

Outline

- Background
  - Impact of AI on science
- Challenges
  - Scientific knowledge is scattered
  - Scientific models can be hard to reuse
- FAIR4AG2 principles
- A case study
  - A plugin for milk spoilage decision-support

AI agents can substantially improve the research workflow

01

Finding the literature

milk spoilage Pseudomonas shelf life

Post-pasteurization contamination of fluid milk by psychrotrophic…
J Dairy Sci · 2019 · Cited by 214

Shelf life of HTST milk: effects of storage temperature…
Food Microbiol · 2017 · Cited by 96

1 2 3 … 27 about 270 results

The search is matched by keyword only. A researcher has to open each abstract, download the PDF, and keep a bibliography manually.

"Find every paper since 2015 that reports a post-pasteurization contamination rate for HTST milk, and tabulate the sampling method used."

- Queried PubMed through a connected tool (MCP): **312** records
- Screened titles and abstracts against the criteria: **41** relevant
- Pulled open full text from PMC; flagged **13** behind paywalls
- Wrote the table: paper · year · n · method · rate · DOI

The search is matched by semantic similarity; the agent searches, screens, and records, and can complete a much more comprehensive search in a much shorter turnaround.

AI agents can substantially improve the research workflow

02

<>

Wrangling and analyzing data

analysis_v7_FINAL.R

```r
df <- read.csv("plates_v3_fixed.csv")
df$log_cfu <- log10(df$cfu + 1)
for (i in 1:nrow(df)) {  # loop by hand
  if (df$day[i] == 7 && df$lot[i] == 9) df$log_cfu[i] <- NA
fit <- nls(log_cfu ~ buchanan(day, mu, lag), ...)
# see also Fig2_v6.R, Fig2_v7.R, Fig2_v7b.R ...
```

Every cleaning rule and model fit written and debugged line by line. Documentation can be a headache as codes accumulate over time.

"Clean the plate-count data, fit the three-phase growth model to each lot, and produce Figure 2 with 95% intervals. Add tests."

`load_plates()` → `clean(rules.yaml)` → `fit_growth()` → `plot_fig2()` → `tests ✓ 14/14`

Describe the pipeline in plain language; the agent writes, runs, and checks the code end to end. New data arrive: rerun with one command.

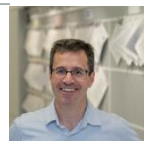

AI companies are investing in AI for science

- Resources are disproportionally abundant in bio science compared to agricultural science

Plugins improved accessibility to scientific database and tools

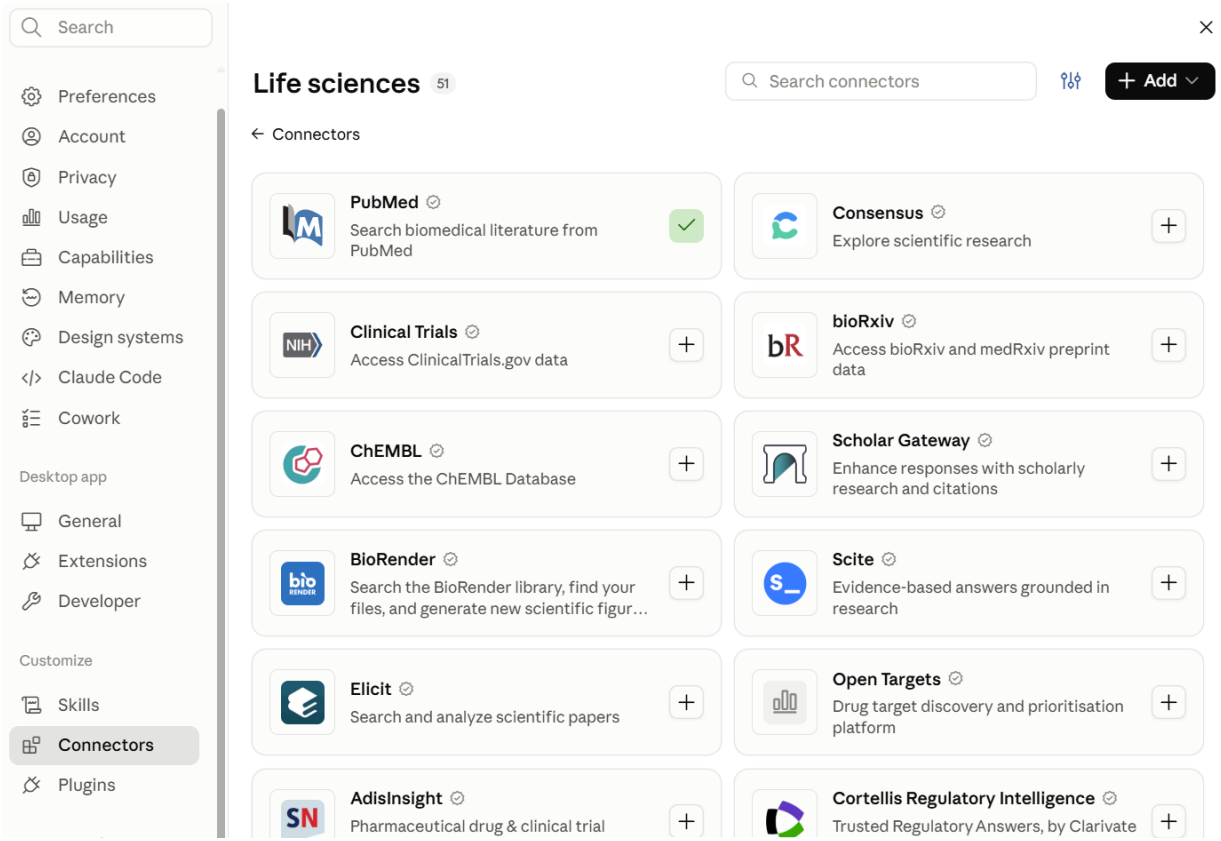

Claude

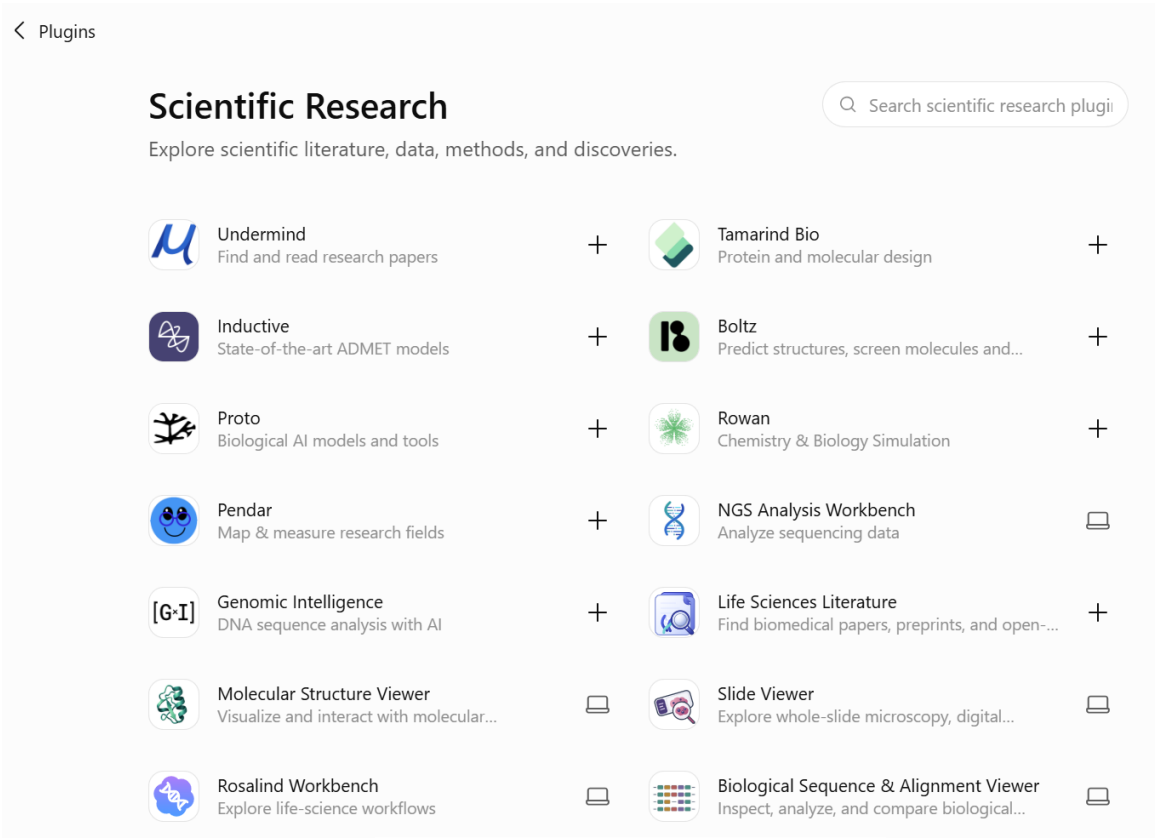

ChatGPT

AI improved scientific writing

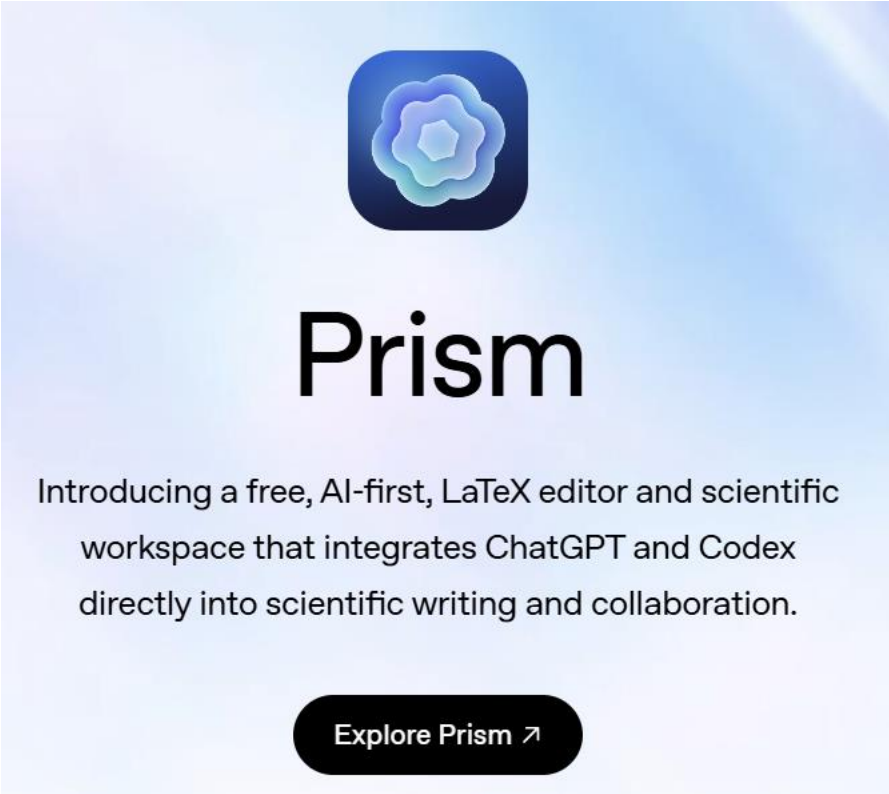

- A free, unified LaTeX (a machine-readable text format) workspace with unlimited collaborators
- AI-assisted proofreading, citations, and literature search
- AI-driven automation for formatting and editing

https://openai.com/prism/

AI workspace for scientific workflow

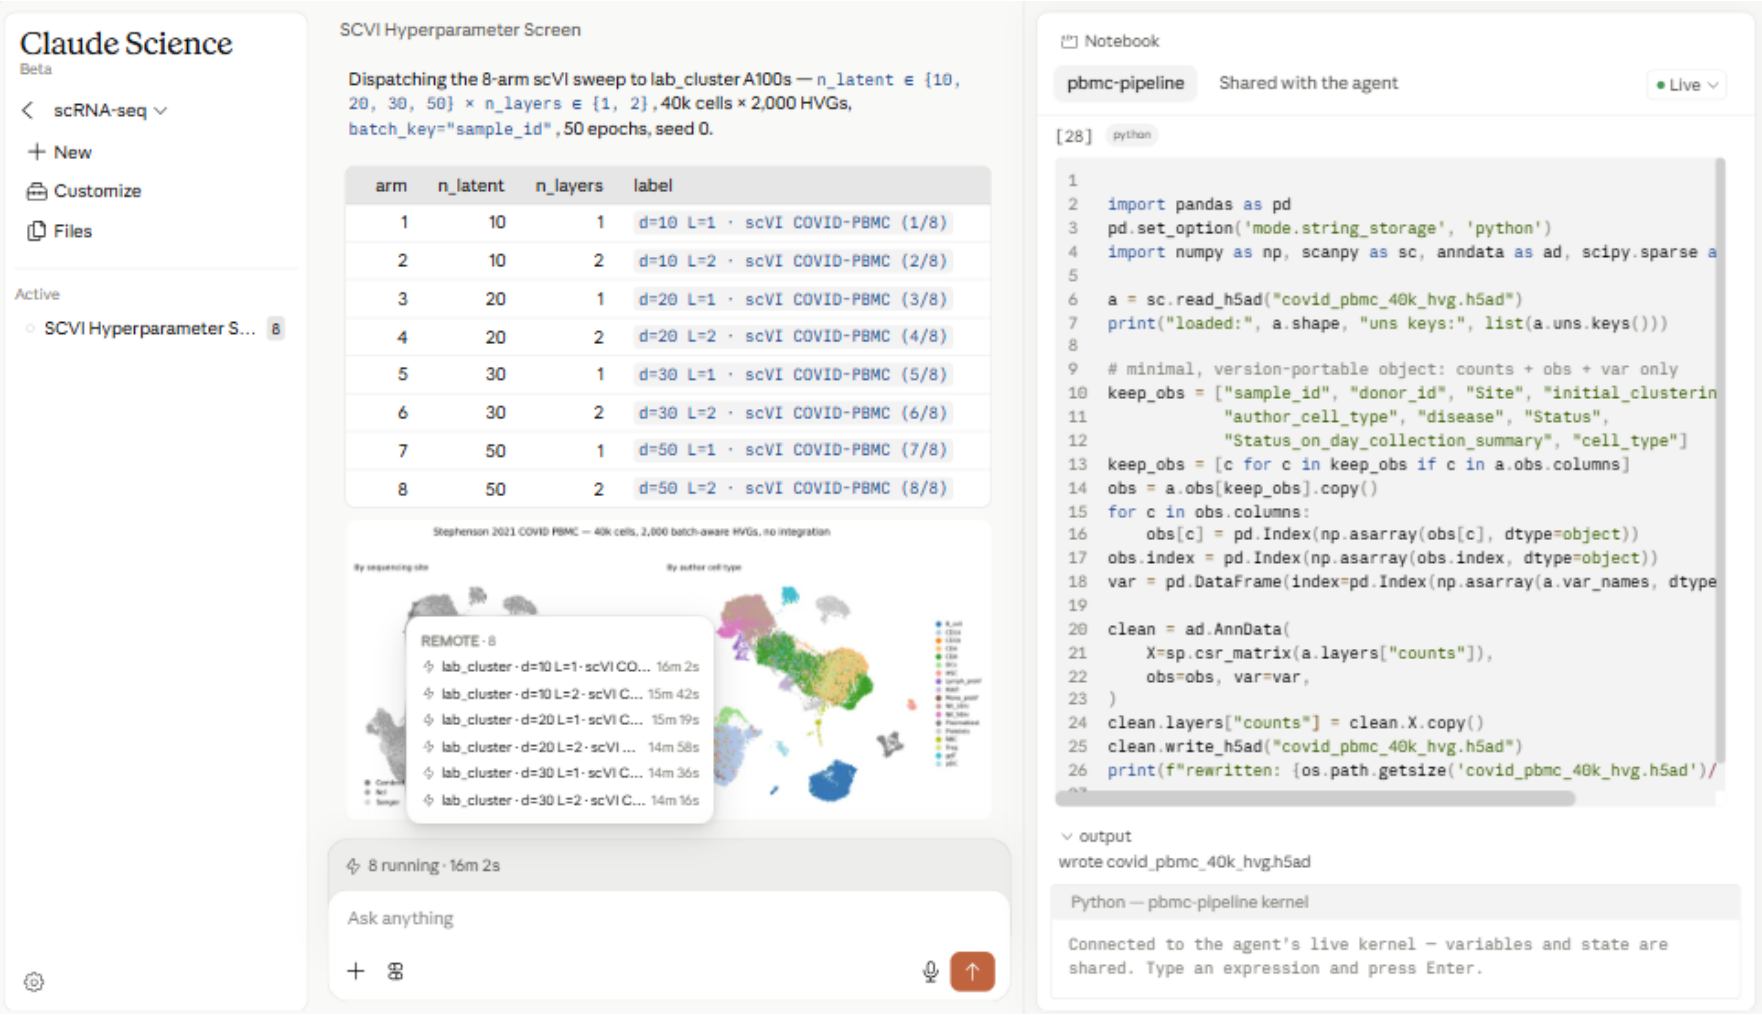

Human and AI have different preferences

HUMAN

AI AGENT

DOCUMENT

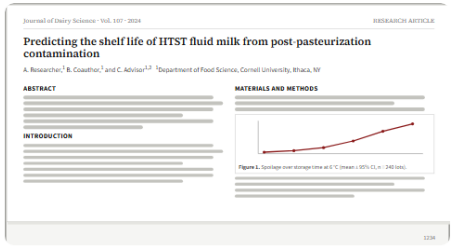

PDF, Word vs **Markdown, JSON, database**

DATA

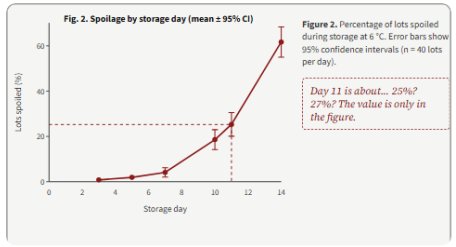

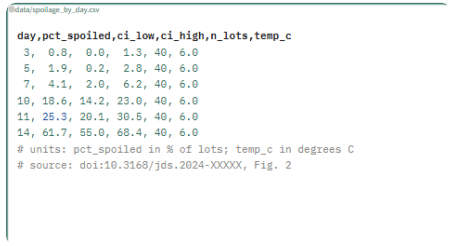

Table, figure vs **CSV**

MODEL

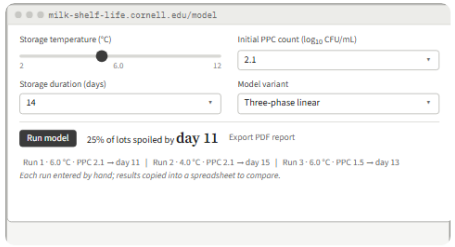

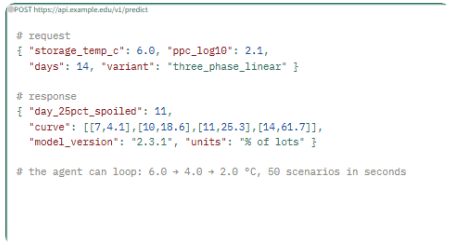

Web interface vs **script, API**

WEB-BASED INFO

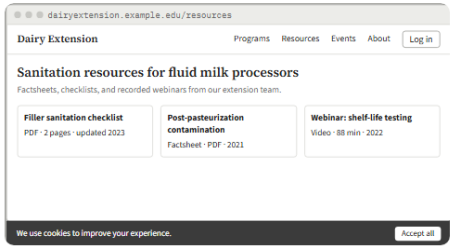

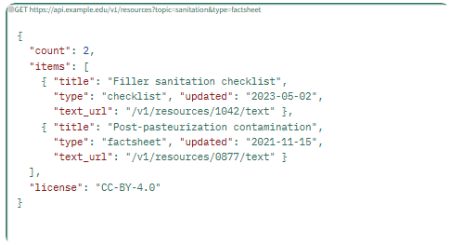

Website vs **API**

SOFTWARE

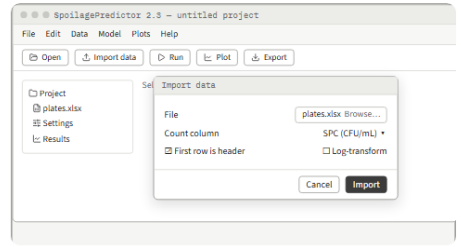

GUI vs **MCP tool**

GUI: graphical user interface (windows, menus, buttons). **API**: application programming interface (a way for software to request data or run a function). **MCP**: Model Context Protocol (a standard way for an AI agent to connect to tools and data). **CSV, JSON, Markdown**: plain-text formats machines read reliably.

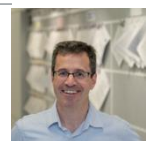

Human and AI have different preferences: documents

HUMAN READS · PDF OR WORD DOCUMENT

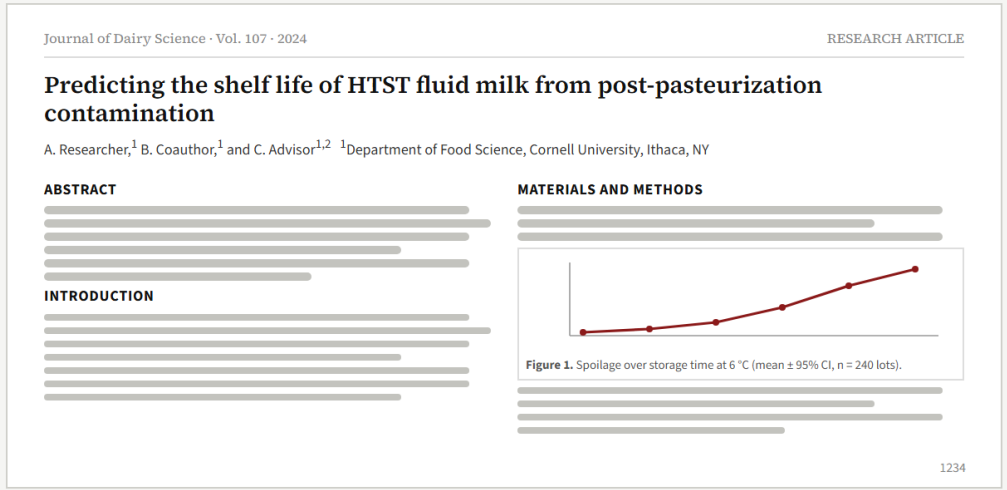
Journal of Dairy Science · Vol. 107 · 2024 RESEARCH ARTICLE

Predicting the shelf life of HTST fluid milk from post-pasteurization contamination

A. Researcher,[1] B. Coauthor,[1] and C. Advisor[1,2] [1]Department of Food Science, Cornell University, Ithaca, NY

ABSTRACT

INTRODUCTION

MATERIALS AND METHODS

Figure 1. Spoilage over storage time at 6 °C (mean ± 95% CI, n = 240 lots).

1234

Layout, typography, and figures for human to read. The parameter values and reported outputs can be somewhere in the text, which can be hard to retrieve.

SAME PAPER

AI READS · MARKDOWN + JSON, OR A QUERYABLE DATABASE

paper.md

```
---
title: Predicting the shelf life of HT
  fluid milk from post-pasteurization
  contamination
doi: 10.3168/jds.2024-XXXXX
model: three-phase linear growth
data: data/plates.csv
---
## Methods
Plate counts (n = 240 lots) were taker
days 1, 7, 10, 14 at 6 °C ...
## Results
Spoilage reached 25% of lots by day 1
```

model.json

```
{
  "model": "three_phase_linear",
  "parameters": {
    "mu_max": 0.42,
    "lag_h":  6.1,
    "unit":   "log10 CFU/mL/day"
  },
  "validation": {
    "rmse": 0.31, "n": 240
  },
  "source": "doi:10.3168/jds.2024-XXXXX
  "version": "2.3.1"
}
```

Structure, fields, and units for AI to read. Every value has a name; the agent reads the number without guessing where it is.

Human and AI have different preferences: data

HUMAN READS · TABLE IN THE TEXT, OR A FIGURE

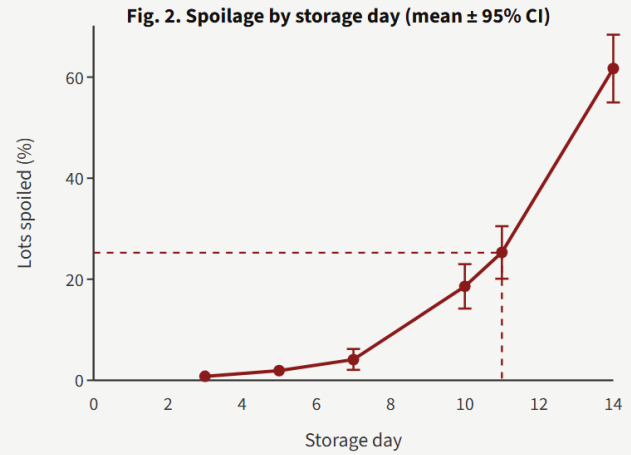

Figure 2. Percentage of lots spoiled during storage at 6 °C. Error bars show 95% confidence intervals (n = 40 lots per day).

Day 11 is about... 25%? 27%? The value is only in the figure.

Data are only presented in the figure. Exact values, intervals, and sample sizes must be read off the axes, if they are reported at all.

SAME DATA

AI READS · CSV (OR ANY TIDY, MACHINE-READABLE TABLE)

data/spoilage_by_day.csv

```
day,pct_spoiled,ci_low,ci_high,n_lots,temp_c
 3,  0.8,  0.0,  1.3, 40, 6.0
 5,  1.9,  0.2,  2.8, 40, 6.0
 7,  4.1,  2.0,  6.2, 40, 6.0
10, 18.6, 14.2, 23.0, 40, 6.0
11, 25.3, 20.1, 30.5, 40, 6.0
14, 61.7, 55.0, 68.4, 40, 6.0
# units: pct_spoiled in % of lots; temp_c in degrees C
# source: doi:10.3168/jds.2024-XXXXX, Fig. 2
```

Data are preserved in a structured format. The agent can plot and refit it, without the need to eyeball the values.

Human and AI have different preferences: models

HUMAN USES · A WEB INTERFACE

milk-shelf-life.cornell.edu/model

Storage temperature (°C)

2 — 6.0 — 12

Initial PPC count (log_{10} CFU/mL)

2.1

Storage duration (days)

14

Model variant

Three-phase linear

Run model

25% of lots spoiled by **day 11**

Export PDF report

RECENT RUNS

Run 1 · 6.0 °C · PPC 2.1 → day 11 | Run 2 · 4.0 °C · PPC 2.1 → day 15 | Run 3 · 6.0 °C · PPC 1.5 → day 13

Each run entered by hand; results copied into a spreadsheet to compare.

Sliders, dropdowns, and a button. The model runs one scenario at a time, by manually changing the input, and the user has to know what each parameter means.

SAME MODEL

AI USES · A SCRIPT OR AN API CALL

POST https://api.example.edu/v1/predict

```
# request
{ "storage_temp_c": 6.0, "ppc_log10": 2.1,
  "days": 14, "variant": "three_phase_linear" }

# response
{ "day_25pct_spoiled": 11,
  "curve": [[7,4.1],[10,18.6],[11,25.3],[14,61.7]],
  "model_version": "2.3.1", "units": "% of lots" }

# the agent can loop: 6.0 → 4.0 → 2.0 °C, 50 scenarios in seconds
```

Named inputs, structured outputs, a version number. The agent can run multiple scenarios and explain the result based on the model developer's instructions.

Human and AI have different preferences: web information

HUMAN BROWSES · A WEBSITE

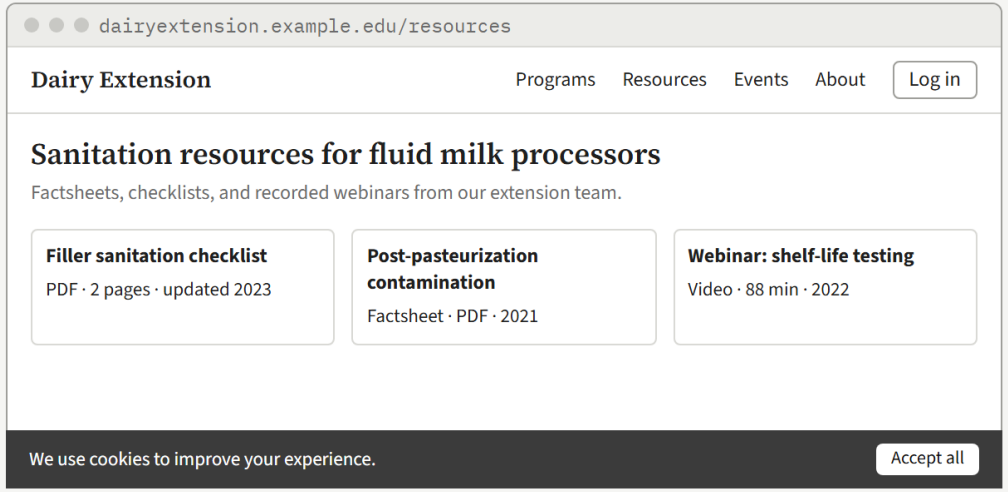

Navigation, layout, and pop-ups. A person clicks through menus; an agent has to use browser extension to read and navigate through bot detection.

SAME CONTENT

AI QUERIES · AN API

GET https://api.example.edu/v1/resources?topic=sanitation&type=factsheet


```
{
  "count": 2,
  "items": [
    { "title": "Filler sanitation checklist",
      "type": "checklist", "updated": "2023-05-02",
      "text_url": "/v1/resources/1042/text" },
    { "title": "Post-pasteurization contamination",
      "type": "factsheet", "updated": "2021-11-15",
      "text_url": "/v1/resources/0877/text" }
  ],
  "license": "CC-BY-4.0"
}
```


A predictable request, a structured answer. Same content, plus the date, the license, and a direct path to the full text.

Human and AI have different preferences: software

HUMAN OPERATES · A GRAPHICAL INTERFACE (GUI)

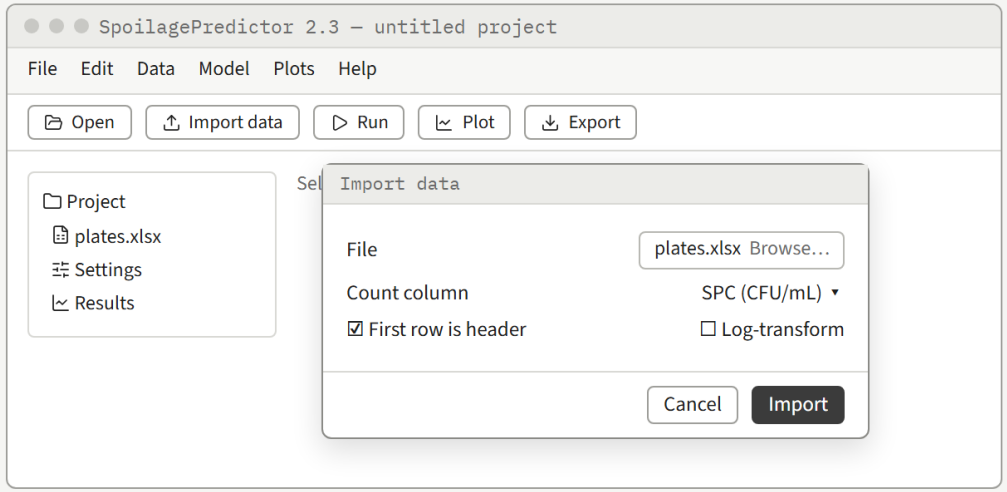

Menus, dialogs, and clicks. Intuitive for a trained person; while AI's computer use ability has been improving, it is less native and much slower to operate.

SAME SOFTWARE

AI CALLS · A TOOL EXPOSED THROUGH MCP

MCP server · tools/list → one tool definition

```
{
  "name": "predict_default",
  "description": "Predict fluid-milk shelf life from plate
      counts. Returns % of lots spoiled by day.",
  "inputSchema": {
    "type": "object",
    "properties": {
      "spc_day1_log10":  { "type": "number" },
      "storage_temp_c":  { "type": "number", "default": 6 }
    },
    "required": ["spc_day1_log10"]
  }
}
```

A name, a description, and a typed input schema. The agent reads what the tool does and calls it directly, no clicking, no training.

Outline

- Background
  - Impact of AI on science
- Challenges
  - Scientific knowledge is scattered
  - Scientific models can be hard to reuse
- FAIR4AG2 principles
- A case study
  - A plugin for milk spoilage decision-support

Research outputs that should be AI-accessible, and why

- **Peer-reviewed papers** so an agent can access the finding and methodology across similar papers, and read the license (which indicates whether agent can redistribute the information)
- **Underlying data** so an agent can re-analyze or reuse for different purposes (e.g., meta-analysis across multiple studies)
    - Possible concerns about us of data to answer questions data were not designed for
- **Statistical and other code** so an analysis can be repeated exactly, audited, and applied to new data.
- **Models** so a prediction can be run on the user's inputs without the hurdle for download and installation.
- **Research outputs from academic or industry organizations** such as extension report, industry guidelines, webinars so that agent can interpret and apply the knowledge properly

Resources are not equally accessible

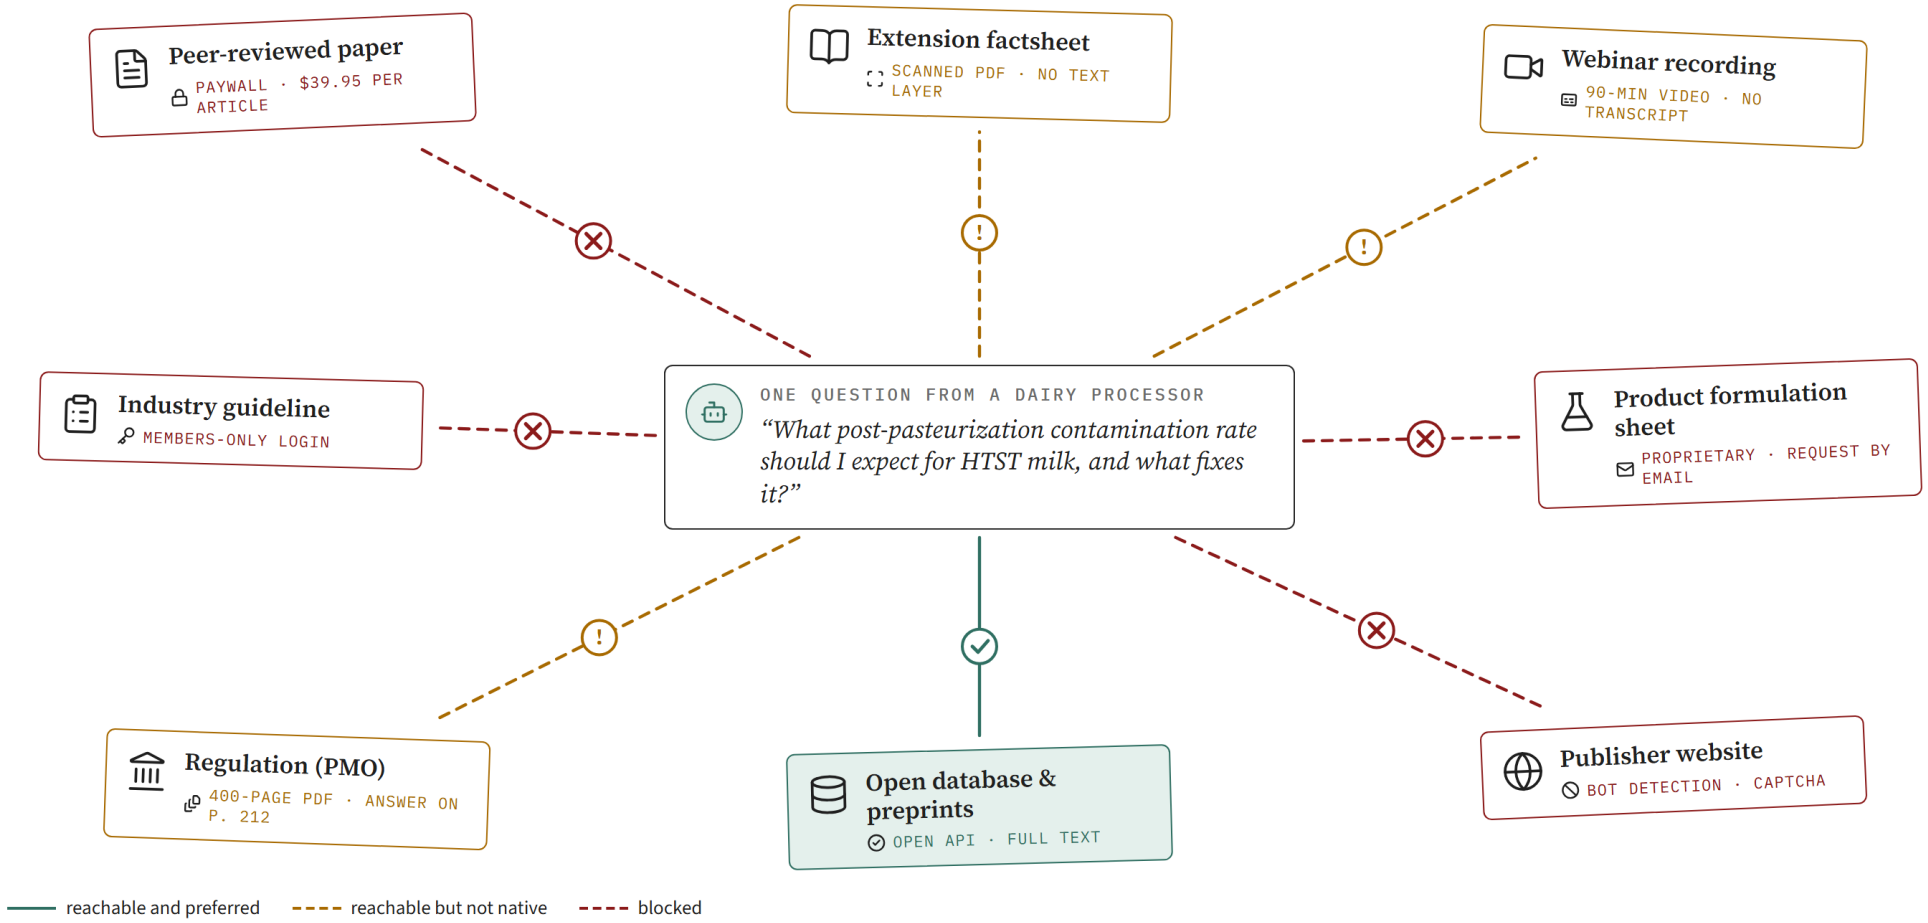

Centralized database in a virtual file-system can help

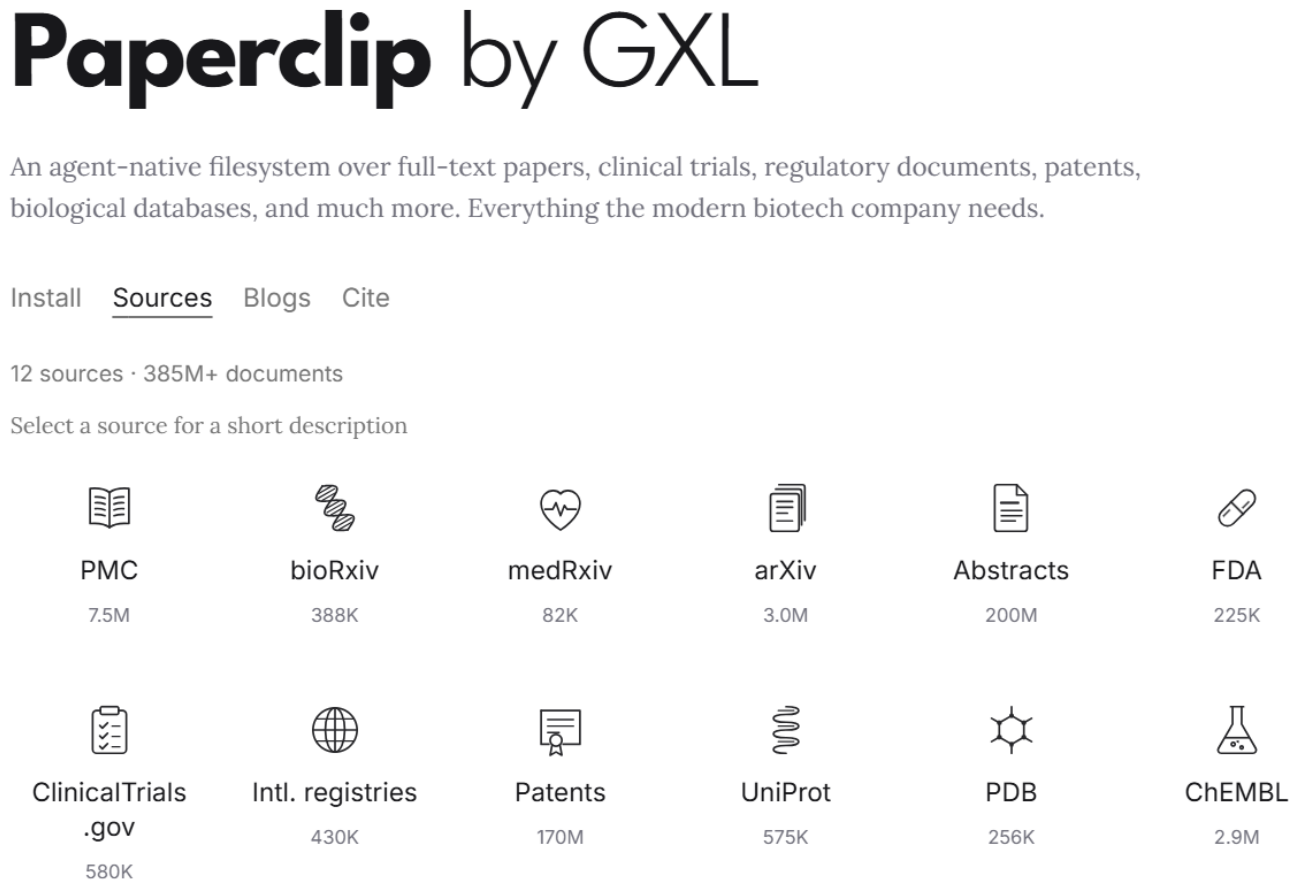

- Agent-native = a system designed primarily for AI agent to operate
- One query searches all relevant databases, faster and more comprehensive than intensive web search
- Stateful = the agent works inside a virtual file system where every paper is a folder, so it can browse, search, and return to what it found instead of starting over with each query

Outline

- Background
  - Impact of AI on science
- Challenges
  - Scientific knowledge is scattered
  - Scientific models can be hard to reuse
- FAIR4AG2 principles
- A case study
  - A plugin for milk spoilage decision-support

Scientific information can be hard to reproduce

VALUES THAT LIVE ONLY IN A FIGURE

Fig. 3. Log reduction by treatment (mean ± SD, n = 3)

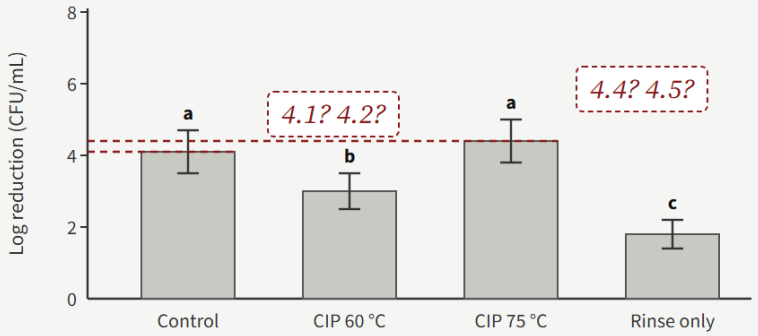

- ✕ No table of values
- ✕ No underlying data
- ✕ No analysis code
- ✕ SD or SE? Legend is silent

Σ AN ANALYSIS THAT CANNOT BE REPEATED

METHODS, STATISTICAL ANALYSIS

"Data were analyzed using a mixed model with treatment as a fixed effect and lot as a random effect (SAS 9.4). Outliers were removed prior to analysis. Means were separated at $P < 0.05$."

- Which covariance structure, which degrees-of-freedom method?
- Which observations were outliers, and by what rule?
- Which multiple-comparison adjustment separated the means?

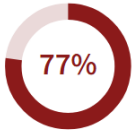

of surveyed **biologists** have tried and **failed to reproduce** another scientist's experiment.

Nature survey of 1,576 researchers; more than 70% across all fields (Baker, 2016, *Nature* 533:452–454).

If a trained researcher struggles to reproduce a result from the paper alone, an AI agent, which has **only** the paper, will struggle more.

Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. Nature 533, 452–454.

Scientific models are hard to reuse

DEAD END 01

Paywalled software

SpoilagePredictor Pro

Your 30-day trial has ended.

$2,400 per seat per year

The model described in the paper is available only inside this application.

Buy license | Quit

The paper is open; the developed model in the paper is not.

DEAD END 02

Repository nowhere to be found

github.com/foodsci-lab/growth-model

404

This repository has been deleted or made private.

"Code is available at github.com/foodsci-lab/growth-model" — Data availability statement, 2019

The link in the paper outlived the researcher who owned the account.

DEAD END 03

"Contact the author"

Mail

To: corresponding.author@university.edu
Subject: Request for model code (J. Dairy Sci. 2019)

Dear Dr. Author, I am trying to apply the growth model from your 2019 paper and would be grateful for the code or parameter files…

Automatic reply: This mailbox is no longer monitored. The account holder left the institution in 2021.

"Available upon request" assumes someone is there to answer.

DEAD END 04

Code drifts from the paper

Which one is "the model"?

PAPER · TABLE 2 (2019)

μ_{max} = 0.42 d^{-1}
λ = 6.1 h
N_{max} = 8.5 log

≠

CODE · V2.3 (2023)

```
mu_max = 0.38
lag_h  = 4.0
n_max  = 8.5
# re-tuned on
# 2022 data
```

The code kept evolving; the paper froze. Results from one cannot be checked against the other.

It is also hard for industry professionals

THE RESEARCHER

Publishes a model to support decisions

A predictive model for post-pasteurization spoilage of HTST fluid milk: application to sanitation decisions

"The model can be used by processors to prioritize interventions and estimate the resulting shelf-life gain."

Peer-reviewed | Validated on 240 lots | Fig. 4: what-if scenarios | Cited 38 times

Intended outcome: a processor runs the model on their own data before choosing an intervention.

THE GAP

- ✕ Paper behind a paywall
- ✕ Written for researchers, not industry users
- ✕ Parameters the reader has to estimate
- ✕ Nothing to run: no tool, no interface
- ✕ No time to read, no one to ask

THE INDUSTRY PROFESSIONAL

Decides from experience

"We have always cleaned the filler every eight hours. Shelf life is fine most weeks. When it is not, we clean more."

Experience | Intuition | Rule of thumb | Trial and error

Actual outcome: the published model is rarely part of the decision.

Executable model repository facilitates model reuse

RAKIP Initiative
risk assessment knowledge integration platform

RAKIP-Web Model Repository

Search Models | Software | Environment | Hazard | Type | 58 models

MODEL	SOFTWARE	ENVIRONMENT	HAZARD	TYPE	EXECUTION TIME	UPLOAD DATE
Dose-response Model of Salmonella with 2D	R	• no information	Salmonella Salmonella Enteritidis Salmonella typhimurium	Dose-response model	12s	2021-11-03 11:05
Dose-response model for Campylobacter jejuni	FSK-Lab	• Any product	Campylobacter jejuni	Dose-response model	59s	2021-11-03 11:08

New tool to improve reusability

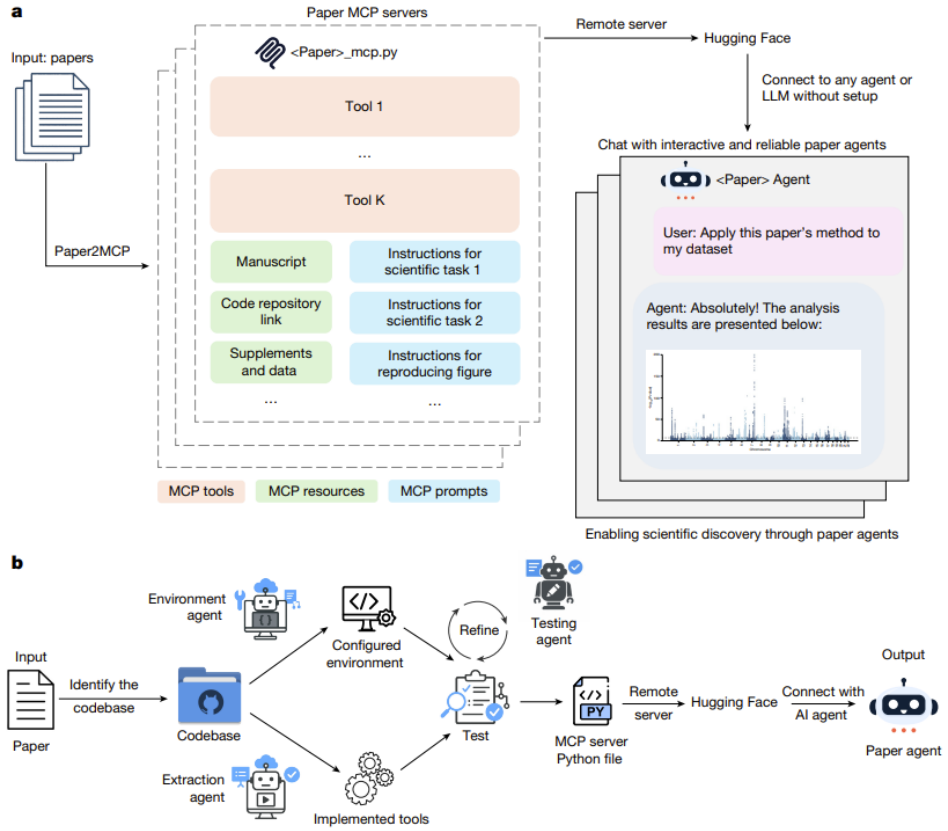

- **Paper2Agent** turns a research paper into an AI agent: it reads the paper and its code, wraps the methods as MCP tools, tests them, and hosts the result so any agent can use it.
- "Reimagining research papers as interactive and reliable AI agents": a scientific outcome is no longer a static document but a conversation with an agent.
- The software tools developed for the paper can be used directly, without rewriting the underlying code.

Miao, J. et al. Paper2Agent: Reimagining research papers as interactive and reliable AI agents. Nature (2026)

Outline

- Background
  - Impact of AI on science
- Challenges
  - Scientific knowledge is scattered
  - Scientific models can be hard to reuse
- FAIR4AG2 principles
- A case study
  - A plugin for milk spoilage decision-support

Existing principles were written for human users

- **FAIR (2016), for research data.** Findable, Accessible, Interoperable, Reusable: persistent identifiers, rich metadata, open protocols, clear licenses. Designed to help both researchers and computational systems find, interpret, and reuse data.
- **FAIR4RS (2022), for research software.** The same principles adapted to code: version-level identifiers, documented dependencies, provenance. Designed to help both researchers and computational systems discover, access, understand, and reuse research software.
- Both make knowledge available to someone who can read a paper and run a script. Neither tells an AI agent how to discover, execute, interpret, or trace a result.

Wilkinson, M. D. et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. Sci Data 3, 160018.
Barker, M. et al. (2022). Introducing the FAIR Principles for research software. Sci Data 9, 622.

We extend the principles to AI agents: FAIR4AG²

2016 · DATA

FAIR

Persistent identifiers, rich metadata, open protocols, clear license. Written for a human at a search box.

2022 · SOFTWARE

FAIR4RS

Version-level identifiers, dependencies accounted for, detailed provenance. Written for a human who runs the code.

2026 · KNOWLEDGE UNITS

FAIR4AG²

Data, models, agent skills and tools. Component-level identifiers, execution interfaces, machine-readable provenance, runtime records.

Findable

From searching for keywords to agents discovering specific skills and tools autonomously.

ADDS

version- and component-level identifiers; registries searchable by agents

Accessible

Beyond downloadable files: standardized interfaces that retrieve data without human assistance.

ADDS

execution interfaces, so models and skills can be called by agents

Interoperable

Version control and metadata for every unit, so an agent can display an explicit chain of thought.

ADDS

machine-readable, version-controlled provenance; alignment metadata for multimodal resources

Reusable

Research software as modular, executable tools an agent can pick up and apply to new problems.

ADDS

runtime provenance logging and machine-readable licensing

Outline

- Background
  - Impact of AI on science
- Challenges
  - Scientific knowledge is scattered
  - Scientific models can be hard to reuse
- FAIR4AG2 principles
- A case study
  - A plugin for milk spoilage decision-support

Case study: a plugin for milk spoilage decision support

- We took a model we developed and turned it into a plugin: an add-on tool that connects to ChatGPT or Claude.
- A dairy processor wants answers to three questions:
  - What is the leading cause of spoilage in our milk?
  - What should we implement to improve milk quality and extend shelf life?
  - Which strategies give the best return on investment?

Why a general AI tool cannot answer these today

- General AI such as ChatGPT gives generic answers. It does not know the processor's own data, for example the microbial counts in their milk.
- The published models that predict milk spoilage and shelf life are not accessible to ChatGPT.
- Even when a model has a web interface, the user needs training to understand the parameters and settings, run it with confidence, and interpret the result correctly.

Solution: package the model as one knowledge unit

- A knowledge unit packaged as a plugin, available to everyone through the OpenAI or Anthropic plugin marketplace.
- Components:
  - An executable model hosted on Hugging Face
  - Tools the agent can run directly, for example predict_default, model_an_intervention, classify_lot_samples
  - A skill document that routes the user's intent to the right tool:
    - Predict the shelf life
    - Compare what-if scenarios
    - Ask for guidance in troubleshooting

A knowledge unit packaged as a plugin

DISTRIBUTION To be listed in the OpenAI and Anthropic plugin marketplaces: anyone with an AI agent can install it in one click.

USER

Dairy QA manager

"Here are ten plate counts from last week. What is spoiling our milk, and what should we fix first?"

Speaks in plain language and their own records. No model parameters, no training.

AI AGENT

ChatGPT, Claude

Reads the skill document, decides which tool fits the question, calls it with the user's data, and explains the result.

General knowledge, plus whatever plugins are installed.

THE PLUGIN · ONE KNOWLEDGE UNIT

Milk spoilage decision support

SKILL DOCUMENT · ROUTES THE USER'S INTENT TO A TOOL

"Predict the shelf life" →	`predict_default` `predict_customize`
"Compare what-if scenarios" →	`model_an_intervention` `model_a_disruption`
"Help me troubleshoot" →	`classify_one_sample` `classify_lot_samples` `read_guidance`

TOOLS · EXECUTABLE INTERFACES THE AGENT CAN CALL DIRECTLY

User inputs (plate counts, temperature, days), structured outputs (percent spoiled by day, dollars, CO_2e), the model version, and an out-of-scope refusal when the question is about pathogens.

EXECUTABLE MODEL

Hosted on Hugging Face

The published spoilage models, packaged to run on request: spoilage classifier, shelf life prediction model, economic and environmental analysis. Versioned, with the paper's DOI attached.

v2.3.1 · doi:10.3168/…

The answer returns in the interpretable language, supplemented by the model's numbers, its version, and the source paper.

The plugin in action

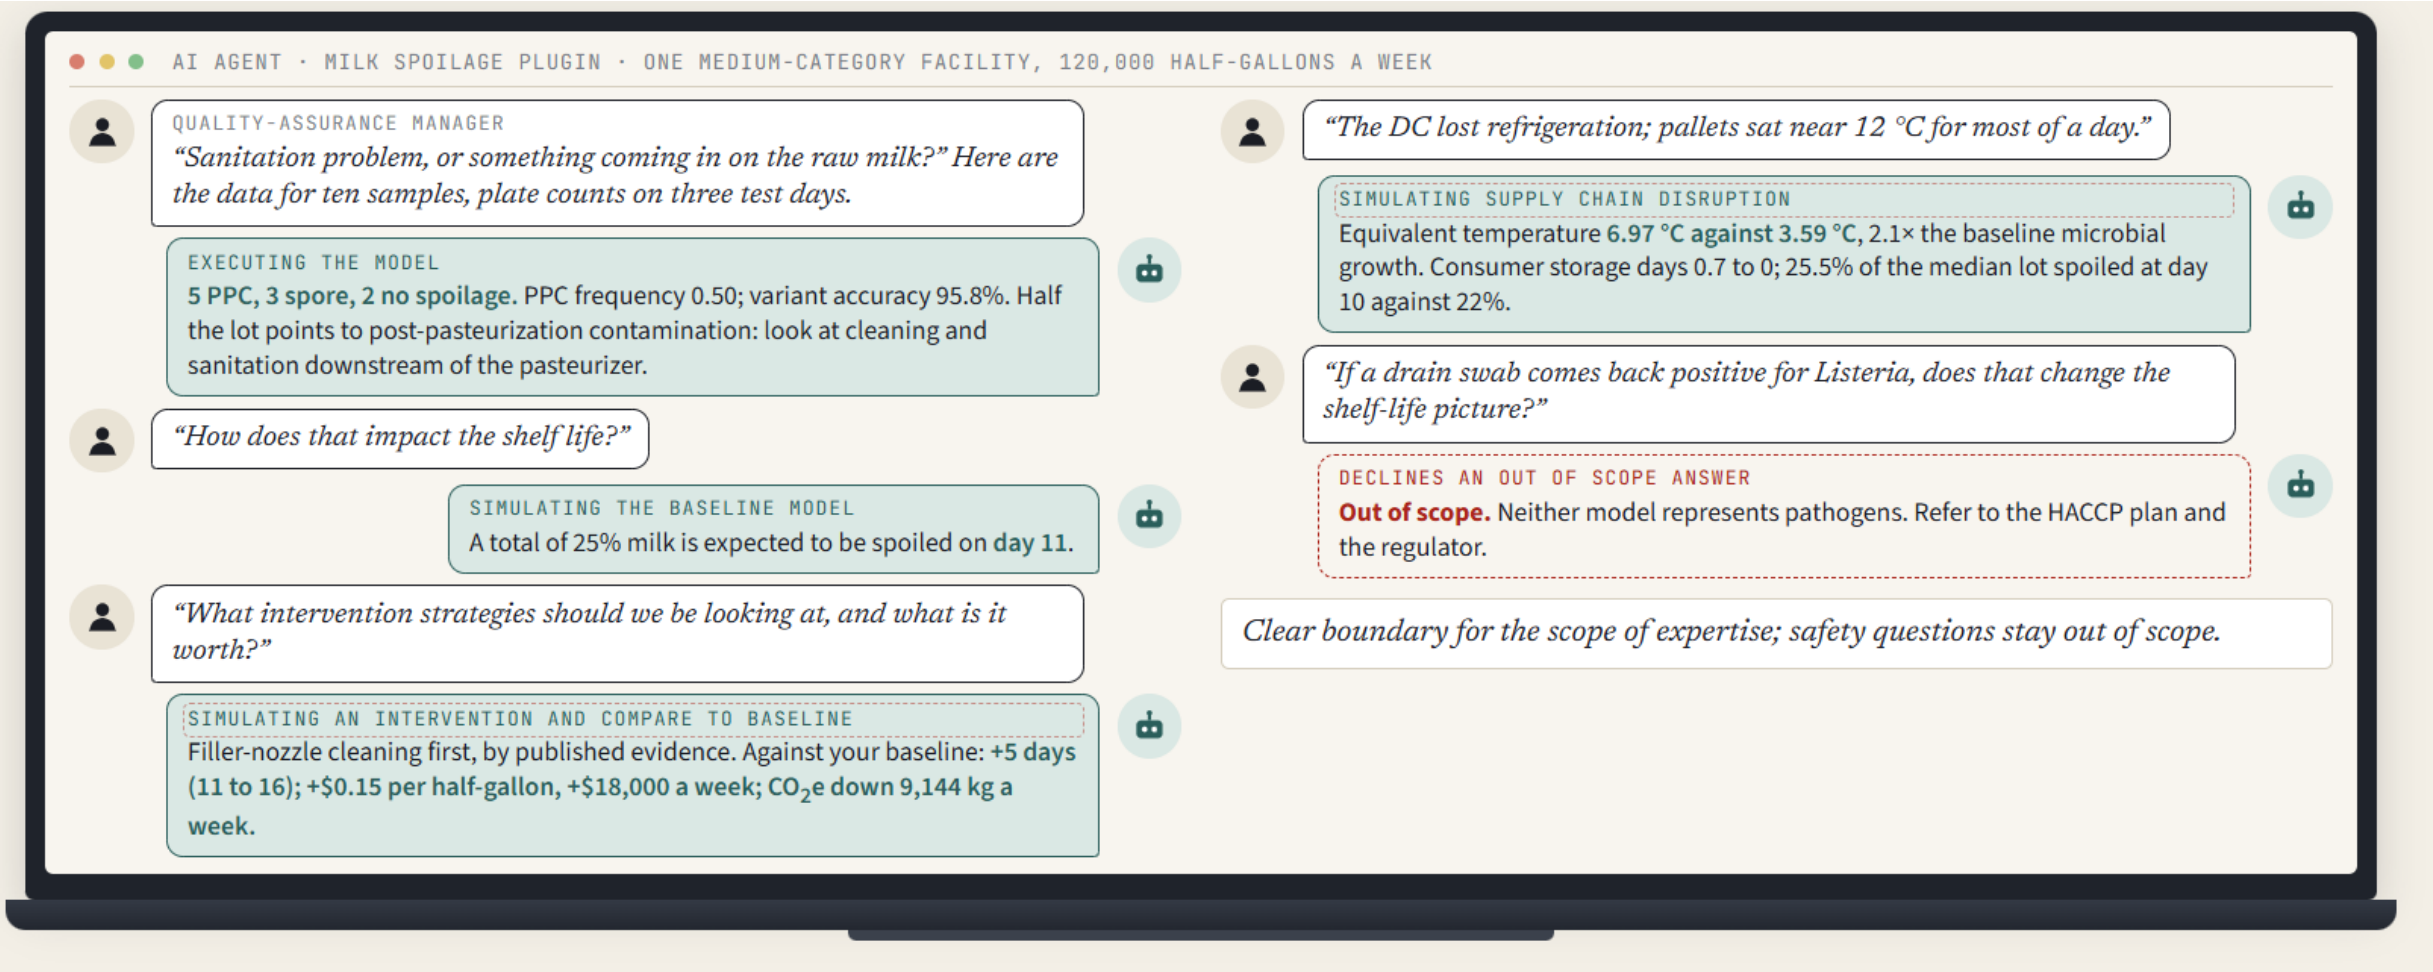

Agent-ready science requires infrastructure

Researchers

Maintain an accessible repository for each published model: code, data, and a skill document, versioned and licensed.

FOR EXAMPLE

A GitHub or Hugging Face repository that an agent can find from the DOI and run without emailing anyone.

Publishers

Adopt standards that make articles machine-accessible: full text and structured metadata through an API, not only a PDF.

FOR EXAMPLE

Parameter tables and data-availability links exposed as structured fields, with a machine-readable license.

Government agencies

Provide MCP servers for the databases, standards, and regulations they already publish, so agents can query them directly.

FOR EXAMPLE

The Pasteurized Milk Ordinance, FoodData Central, and NASS statistics as queryable endpoints, not 400-page PDFs.

Funding agencies

Financially support the long-term maintenance of executable models beyond the grant that produced them.

FOR EXAMPLE

Hosting, version updates, and a named maintainer for five years after publication, written into the award.

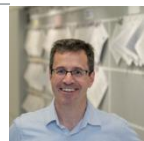

Risks of agent-actionability

- An unsuitable model can be applied quickly and repeatedly.
- Invalid inputs can produce convincing explanations.
- Agents may continue using outdated model versions.
- Data may be used to address questions the data were not intended to answer
- Model developers and data curators may not be attributed with credit properly.
- Unequal infrastructure may widen existing research disparities.
- **High-stakes decisions still need a person to check the answer.**

Barriers: why researchers hesitate to make work agent-ready

- **Fear of exposure.** Once code and data can be run by anyone, it is easy for others to show that a result does not hold.
- **Unpaid work.** Packaging and maintaining a model is not funded by the grant and not required by the journal.
- **Skills.** Many labs have no one whose job is to build and maintain software.
- **What helps:** treat reproducibility as a design goal from the first day of a project, share infrastructure, give credit for reusable outputs, and use AI agent (e.g., Claude code) effectively to build accessible knowledge unit

Future-proofing research

- A research paper can document what we learned.
- But an agent-ready research output can keep that learning usable by the entire scientific community, and by the industry it was meant to serve.
- **Which research output in your field should become agent-ready first?**


PERSPECTIVE article

Front. Artif. Intell., 05 August 2026

Sec. AI in Food, Agriculture and Water

Volume 9 - 2026 | https://doi.org/10.3389/frai.2026.1877344

Future-proofing agricultural research: FAIR principles for agriculture AI agents (FAIR4AG²)