PURSUIT OF ZΩH — THE WEEKLY

This Week in

AI & Machine Learning

NO. 007 • 15 NOVEMBER 2026 • NEWS OF 8 NOV – 15 NOV 2026 • 3 STORIES

THIS WEEK

Machines that pitch ideas, and experts who grade them

An AI pitches over 4,400 research ideas to 100-plus lab heads, a climate emulator learns more from one odd scenario than from six standard ones, and a new score shows when merging cell data pays off.

IN THIS ISSUE

01 AI FOR SCIENCE
An AI pitched over 4,400 research ideas; experts were lukewarm

02 MACHINE LEARNING FOR CLIMATE
One strange training scenario beats six realistic ones

03 COMPUTATIONAL BIOLOGY
When merging cell data adds something new, and when it doesn't

01

AI FOR SCIENCE

An AI pitched over 4,400 research ideas; experts were lukewarm

Xuemei Gu and Mario Krenn, at the Max Planck Institute for the Science of Light, built SciMuse, which pairs a knowledge graph of 58 million papers with a large language model to propose ideas tailored to one scientist.[1] More than 100 research group leaders scored over 4,400 such ideas for interest.[1] The mean was 2.40 out of 5, and 24.9% earned a 4 or 5.[1] Graph-picked concept pairs did no better than a GPT baseline given only paper titles.[1]

WHY IT MATTERS

Interest is hard to measure, and this rating set let the authors predict it, both with a neural network on graph features and with a language model ranking ideas zero-shot.[1]

CAVEAT *Pairs predicted to draw many citations leaned slightly toward lower interest than random pairs, a weak 1.94σ tendency.*[1]

[1] doi:10.1073/pnas.2611349123

24.9%

of the 4,400-plus AI-generated ideas that research group leaders rated 4 or 5 on a 5-point scale

SOURCE 1

ON THE ROAD?

Listen to this issue

6-minute episode, read aloud: the same three stories, written for the ear. Stream it, or download the MP3, from the AI & Machine Learning page.

pursuit-of-zoe.com/pages/ai-and-machine-learning#zoh-weekly

02

MACHINE LEARNING FOR CLIMATE

One strange training scenario beats six realistic ones

Climate emulators are fast AI stand-ins for Earth system models, which can need months on supercomputers per scenario.[3] An MIT-led team reports in PNAS that the sameness of the usual training pathways caps how well emulators generalize.[2,3] Using a differentiable simple climate model, they reshaped the training scenario step by step to cut emulator error.[2] One optimized scenario beat six standard ScenarioMIP pathways as training data, despite the smaller dataset.[2]

WHY IT MATTERS

A modeler would run the costly full model once rather than six or eight times, and the emulator separated greenhouse-gas and aerosol effects without dedicated single-forcing runs.[2,3]

CAVEAT *The headline comparison uses a simple climate model; the paper's step up is to an intermediate-complexity model, not a full Earth system model.*[2]

[2] doi:10.1073/pnas.2620731123 · [3] news.mit.edu

MECHANISM

Start from an arbitrary emissions scenario

↓

Simple climate model scores emulator error

↓

Gradients reshape the training scenario

↓

Repeat until the error hits a threshold

↓

Train on the wiggly, unrealistic result

03

COMPUTATIONAL BIOLOGY

When merging cell data adds something new, and when it doesn't

Cell biology has strong foundation models for single data types, such as gene expression or tissue images, but few large paired datasets.[4] Writing in Nature Machine Intelligence, Till Richter and colleagues introduce the Synergistic Information Score, which measures information that appears only when data types are combined.[4] Standard alignment objectives with frozen encoders end up detecting linear redundancy, and across ten fusion methods on spatial transcriptomics data, redundancy-dominated tasks were well served by single-type baselines.[4]

WHY IT MATTERS

Complex definitions of tissue niches did gain from synergy-aware fusion, so the score helps show when a multimodal model is worth building.[4]

CAVEAT *In the authors' scaling analysis, fine-tuning the strongest single-type model was the most sample-efficient strategy for standard tasks.*[4]

[4] doi:10.1038/s42256-026-01328-3

MECHANISM

Frozen experts: gene expression, images

↓

A learned interface fuses their outputs

↓

Score isolates cross-modal-only information

↓

Mostly redundant: one expert suffices

↓

Real synergy: fusion pays off

WORDS THIS WEEK

Knowledge graph
A network whose nodes are concepts and whose links record how they connect, here built from which ideas appear together across millions of papers.

Climate emulator
A fast statistical or machine-learning model trained to mimic a slow, physics-based climate model, so many scenarios can be tested in seconds.

Foundation model
A large model pretrained on broad data, here one type of cell measurement, that can then be adapted to many specific tasks.

Spatial transcriptomics
Methods that measure which genes are active in a tissue section while keeping track of where in the tissue each reading came from.

Differentiable model
A model written so a computer can calculate how its output changes when any input is nudged, which lets an optimizer adjust the inputs automatically.

SOURCES

1. PAPER X. Gu & M. Krenn, "Generating interesting scientific ideas using knowledge graphs and LLMs: Evaluations with 100 research group leaders," Proceedings of the National Academy of Sciences (9 Nov 2026). doi:10.1073/pnas.2611349123
2. PAPER C. B. Womack et al., "Optimal scenario design for climate emulation," Proceedings of the National Academy of Sciences (9 Nov 2026). doi:10.1073/pnas.2620731123
3. RELEASE "New research could lead to faster, trustworthy AI models for climate prediction," MIT News (12 Nov 2026). https://news.mit.edu/2026/new-research-could-lead-faster-trustworthy-ai-models-climate-prediction-1112
4. PAPER T. Richter et al., "Multimodal cell foundation models require synergistic integration beyond alignment," Nature Machine Intelligence (9 Nov 2026). doi:10.1038/s42256-026-01328-3

How this is made. Each story comes from the past seven days of journal, agency and society releases and is written fresh for curious minds, with a source for every claim. Superscripts refer to the list above. Spotted an error? Tell us through the contact page; corrections run in the next issue.

Choose what we feature next. Signed-in customers vote every other week at pursuit-of-zoe.com/pages/vote.