


| Behind the AI

Junseon Yoo | Founder & CEO, Pluto Labs

Welcome

Who is speaking



**I build these tools.
I also do the research.**

mechanical-engineering researcher, still writing CS papers

Welcome

Who is speaking

How perception of AI writing papers changed

2018

I estimated my observation data with deep learning and used it as my main evidence. Technically, I generated my data with AI.

No one questioned the accuracy.

No one worried about AI writing the paper.


TODAY

everyone worries about AI writing papers

The Promise

It is a way to evaluate any AI research tool

Solving the Gap



What I want to give you

is a clearer picture of the research environment:

— how research works and how researchers actually respond to AI

Motivation: Researchers do not use AI

Researchers still reach for Google Scholar, not ChatGPT

AI is everywhere — every piece of software claims AI features, even Gmail offers to write your emails


From our tool sessions

Every time I ask the room
— “do you use AI in your actual
research?”

Not one hand

Search Interest Trends

Source: Google Trends – Last 5 years worldwide Trend



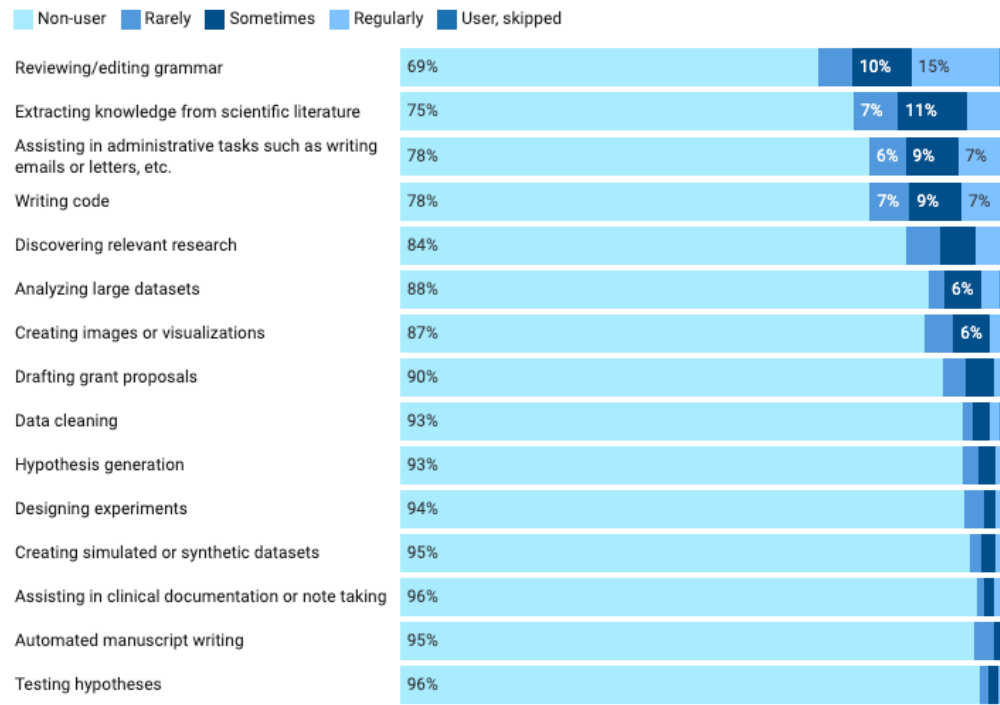
Researchers do not think using services like ChatGPT is superior to existing research methods.

Motivation: Researchers do not use AI

Researchers use AI for grammar, not for research


How researchers use AI, and how often

How often do you use generative AI for the following tasks?



Source: Ruediger, Dylan, et al. "Adoption of Generative AI by Academic Biomedical Researchers." Ithaka S+R. Ithaka S+R. 17 October 2024. Web. 30 September 2025. <https://doi.org/10.18665/sr.321415>

Purpose of AI using in research



Source: Chugunova, Marina, et al. "Who Uses AI in Research, and for What? Large-Scale Survey Evidence from Germany." Research Policy, vol. 55, no. 2. Elsevier. 19 May 2025. Web. 22 June 2026. <https://doi.org/10.1016/j.respol.2025.105381>

Key Question

Why don't researchers use AI in their research in the first place?

01 What an LLM is and why it breaks

Questions From Librarians

01

Biggest warning everyone should know about AI.

University of North Dakota

02

How does one know when there is a risk of hallucinations vs not?

Towson University

03

Cautions needed in using tool--- like driving a car without any instruction--need to know where the "brakes" and "steering wheel" are!

University of Michigan

04

Mitigating risk from using AI; confirming accuracy of info derived from AI and analysis from AI

Laredo Public University

To answer any of these — we need to start one level down : how an LLM actually works

Pillar 1: What an LLM is and Why it breaks

LLMs predict the next word. They do not understand.

Next token prediction

“The capital of Korea is _____”



What it is doing

Probability matching
— not understanding.


It is not retrieving a fact.
It is not reasoning to a conclusion.

Pillar 1: What an LLM is and Why it breaks

LLMs predict the next word. They do not understand.

Vector embedding: For an LLM to predict what comes next, every word/token has to become “Vector” – a set of numbers.

The closer two vectors sit, the more related the model thinks the words are



Ex) What we usually say:

Eat an apple / Eat a banana

→ sits closer and the model decides apple and banana must be similar things.

The precision that matters to a researcher is the precision the model compresses away.

3D structured porous
ceramic surface

Cylindrical
metal surface

collapsed

“3D structured surface”

Pillar 1: What an LLM is and Why it breaks

Why LLM makes things up

LLM never says “I don’t know”

Reason 1:

**It answers based on
“Training Data”**

The model never learned that
“I don’t know” is a normal
response.

Reason 2:

It takes a Guess

There's no penalty for guessing.
So the model always guesses,
even when it doesn't know.

*OpenAI published a whole paper on this,
and it even made it into Nature.*



**AI tries, by probability, to
generate something that
looks like an answer.**

Ex) “What smartphone did Abraham
Lincoln use?”

Instead of saying “there were no
smartphones back then,”
the model would answer with whatever
is closest to “what is a smartphone”

Pillar 1: What an LLM is and Why it breaks

Built-in limits: fabricated references, weak reasoning, stale knowledge



Fabricated references

LLMs invent citations that look completely real — but the paper doesn't exist.

This is where "AI slop" in your databases comes from.



Weak non-linguistic reasoning

LLMs predict text, not logic.

Earlier models couldn't even tell that $13.8 > 13.11$.

Newer reasoning models are better — but the underlying mechanism hasn't changed. It's still text prediction, just more refined.



Confirmation bias

Agrees with whatever direction you push it toward.



Stale knowledge

LLMs only know what existed when they were trained.

Anything published after that point simply doesn't exist to the model.

02 How AI tools find information

Questions From Librarians

01

How to categorize and differentiate AI features appearing in library resources — natural language search, content summaries, and generative AI (ChatGPT-style) — and how concerned they should be about each

Westfield State University

02

How AI technically searches for information and how it could be plugged into our systems

university of Vermont

03

I'd like to know more about semantic search and the "meaning space" framework

University of Virginia

04


how databases are using AI and how that affects the way users use them

Kennesaw State University

Pillar 2: How AI tools find information

Do not chase the keywords, Focus on the Questions

When you try to understand how AI tools go beyond LLM limits to find information —and the keywords explode.



Ask 2 Questions

1. Search once or refine?

Does it search once and stop — or does it take multiple passes, refining the query each time?

2. What kind of search does it run?

more precise as you go.

1. Keyword matching
2. structured database queries
3. graph traversal

Pillar 2: How AI tools find information

How we solved it: a citation graph, redrawn for every question

Draw Citation Graph within a specific field,
or within a specific country, almost in real time

Citation Network

Energy Storage

Solar thermal systems

silicon efficiency limits



Cluster papers by infield citations
→ Surface the agendas within a field

Agentic Workflow:

we wrapped citation network
inside an agentic workflow

AI: Plan
"Which papers to target?"

AI: Search
paper based on citation network

AI: Check
If results are sufficient

(If not) AI: Refine
Adjust query → re-search

Redrawing the graph
every time

We can tell you:

how many times that paper was
cited within a specific field, or
within a specific country, almost
in real time

Others:

a paper's total citation count.
—A total count tells you little
across fields

Pillar 2: How AI tools find information

So why citations?

Why is a citation such a special signal? Look at how a researcher actually uses one.

A citation is load-bearing

Citing a paper isn't saying "this is related."
It means "I built my work on top of this." If that paper collapses, yours shakes too.

Repeated citations $\neq$ popularity

When independent groups keep building on the same work — over years — it means the work held. Refuted research loses its place over time.

The field's own record

Citation graph is, in the end, the field's own record of what it stands on — one that corrects itself.

One citation is not enough

A citation is not a certificate of quality. Influential papers can still turn out to be wrong — and a single citation is never a replication.

What we actually measure
:Not raw citation count — but whether a paper holds its place over time.

03 Where AI actually fits

Questions From Librarians

01

How will AI transform the way research is conducted? What do librarians need to know to prepare for the future?

University of Richmond

02

I want to know how to keep the library relevant, through new tools, or by creating our own.
I am curious about how this will advance federated search and what the implications are.

West Coast University

03

Are there clean-up projects that could be more successful using AI? Are there practical ways I should be using AI in my day-to-day? What sorts of things should we use as a rubric when evaluating AI? How can we really understand the environmental cost of using AI?

University of Chicago

Pillar 3: Where AI actually fits

One more word to put down: “agent”

“AI agent” has become a buzzword — everyone says it, everyone means something different.

Strip the label, Start from what researchers actually do.

agent. agentic. AI agent. agent...

Q1 How can AI take a piece of research work?

Start from what researchers actually do
and then ask Which parts of their research
workflow AI can handle

Q2 Does AI do it well?

"can do it" and "does it well" are
two very different things.
The question is whether AI does it well.

This same lens works for your own world too: Metadata, cataloging, reporting ...

Pillar 3: Where AI actually fits

Research Workflow

 = *What we cover*

1. Knowing

Research question

Search papers

Track trends

Judge quality

Map structure

Spot gaps

2. Doing

Hypothesis

Novelty check

Experiment design

Run / collect data

Analyze data

Interpret / verify

Key findings

3. Telling

Write paper

Make figures

Peer-review reply

Pillar 3: Where AI actually fits

The researcher's workflow – and where AI actually fits



What actually enrages researchers is **when KNOWING fails and DOING breaks**:
— months on a question someone already answered, or a hypothesis that never had a chance.

Doing KNOWING well prevents the most painful waste in research.

Where we sit: **KNOWING + the hypothesis-checking part of DOING.**

Two questions to take to any AI research tool

“Everyone says their tool is the best.” These two cut through all of it — keep them in your pocket.

1 How does it find information?

- Surface similarity, or the structure of how a field is built?
- What is the corpus, and how is it curated?
- Can every claim trace back to a real source?

2 Does it know its place?

- Which part of the workflow does it actually do — and how well?
- “Can do it” is not “does it well.”
- A tool that claims to do everything is telling you it does not understand the work.

The same lens works for your own work, too — metadata, cataloging, reporting.

Additional Question

Can AI write a Research Paper?

Why a paper that meets a journal's bar is, right now, close to impossible for AI

Questions from Librarians:

- How to spot AI written academic papers
- How can we expect AI to affect scientific publishing? At what point will AI-generated knowledge pass peer review checkpoints?

— Librarian, University of Houston

— Librarian, Texas Medical Center

The Best Case so far : Sakana AI (~\$400M funded) got one AI-written paper past peer review.

But look at the details:

- A workshop, not a main track (ICLR)
- Score: 6.33 — barely over the line
- Run as an agreed experiment with organizers
- Withdrawn before publication
- In CS only — where experiments are code, not physical

One of the best-funded AI labs got an AI paper over the bar exactly once — with help, barely, and pulled it.

AI can help you know. It does not do the research for you

Three things to take with you

What LLM is : A next-word predictor.

Powerful — but on its own,
it guesses.

And it gets worse exactly where
research lives: the precise, the new,
the edge cases

LLM is not a research tool.

Why RAG exists

Every serious tool wraps the model
to find real information, not guess.

The question isn't "does it use AI"
— it's how does it find information?

See if it surfaces similarity,
or the actual structure of a field.

How to evaluate : two Questions

How does it find information
— and can every claim trace to a
real source?

Does it know its place —
what part of the work does it do,
and how well?

Do not fear it. Do not buy the hype

In a world full of fluent, confident machines, the people who ask “what is this built on, and can I verify it” are exactly who the academy needs most. That is you.

Q&A

We'd love to answer your questions.

Thank You

A good tool delivers its true value only on the foundation of a good relationship.
I hope today is the start of one.

PRESENTER

Junseon Yoo
CEO

Questions are welcome

junseon@pluto.im



Connect on LinkedIn



scinapse.io
ai.scinapse.io

An AI partner that thinks like a researcher.