

Use AI appropriately as you grow at Purdue.

Match the help to what you are learning, what you can check, and what is allowed.

BEFORE EACH USE OF AI, ASK

- 1** **PURPOSE**
What is the purpose?
Name what you want to learn or accomplish.
- 2** **PERMISSION**
What is allowed?
Check the rules for assistance and sharing data.
- 3** **REASONING**
What must I do myself?
Protect the reasoning you need to learn.
- 4** **CHECKING**
How will I check it?
Verify the result and explain your contribution.

Your use should change as your understanding grows.

● Early coursework → ● Later coursework → ● Projects & research → ● Your career

Learn to use AI as part of your Purdue education.

1 What capable AI can do now

Provider reports: GPT-6 Astra [1] • Claude Fable [2]



2

What surveyed employers expect

35%

of their entry-level jobs required AI skills. [3]

Also valued: communication, teamwork, critical thinking. [3]

NACE Spring 2026 survey: 185 respondents overall; counts varied by question. [3]

3 What employment data show

Stanford: employment for young workers lagged in jobs with more tasks that AI can do. [4]

Yale: broader data show no clear overall effect of AI on jobs. [5]

Different populations and methods. Neither establishes that AI caused a change or predicts your job market at graduation.

Learn to use capable tools while building knowledge, communication, and judgment.

Provider reports, employer surveys, and payroll data answer different questions.

Sources: [1] OpenAI, Astra (2026) ↗ [2] Anthropic, Fable (2026) ↗ [3] NACE (2026), pp. 10, 24, 40 ↗ [4] Brynjolfsson et al. (2026) ↗ [5] Yale Budget Lab (2026) ↗

Employer expectations and employment trends answer different questions.


Human teams and AI analysts split on the same data.

PUBLISHED STUDY • BERTRAN, FOGLIATO & WU • PNAS, JULY 2026

[1]

Same soccer data and question: are referees more likely to give red cards to darker-skinned players?

The data: 2,053 players in the top men’s leagues of England, Germany, France, and Spain in the 2012–13 season, each paired with all the referees they played under in their careers (146,028 pairs); skin tone rated from photos by two raters. [3]



- Each team and AI analyst chose its own analysis; “found significant bias” means that analysis gave $p < .05$ for more red cards to darker-skinned players. [2]
- 4 Caution: the agents had no internet access, but these findings are widely published, so a model may have recalled them from its training data.

Treat one analysis as one path, not the answer.
Defensible choices can flip a verdict, for people and for AI. Before trusting a result, ask which choices drive it.


Check what drives an AI analysis’s verdict.

PUBLISHED STUDY • BERTRAN, FOGLIATO & WU • PNAS, JULY 2026

[1]

1 Same data, same question: each dot is one AI analysis [2]

- 122 said “supported”
- 566 said “not supported”
- no issues flagged by the AI auditor
- minor issues flagged



THE QUESTION

“Do people who watch more TV news show a tighter link between symbolic ideology and policy positions?”

WHAT EACH DOT SHOWS

Its estimated effect of 1 SD more TV news on the ideology–policy gap, in standard deviations

EACH AI ANALYST CHOSE

covariates, regression method, standard errors, and pooling of election years

2 FRAMING STEERS THE VERDICT

13% → 43%

of runs on this question said “supported”: a prompt calling the hypothesis implausible vs. one pushing the agent to find support.

3 THE AUDIT CAUGHT MUCH OF IT

45% flagged

566 of 1,270 runs had major issues or were noncompliant. Among acceptable runs, the find-support prompt’s rate fell to 20%.

Don’t lead the AI: a leading prompt can bias the result, like a leading survey question.


Ask it to test the claim, not to find support. Try other reasonable analyses, and report all results with your prompt.

Sources: [1] Bertran et al., PNAS (2026) [2] Data: released runs (GitHub)

Figure and card numbers from the released runs [2]. Dots: 688 ANES analyses rated compliant (solid) or minor issues (hollow) by the AI auditor; 5 estimates beyond each axis end are drawn at the ends.

Judge AI’s help by measurement, not impressions.

PUBLISHED RANDOMIZED TRIALS • METR • REDRAWN FROM THE 2026 UPDATE’S FIGURE



Judge AI’s help by measurement, not by how it feels.

Then check who and what was measured, and when: what a study leaves out can bias its result.

Sources: [1] METR (2025) ↗ [2] METR update (2026) ↗

Redrawn from METR’s 2026 figure [2], adding developers’ estimates from [1]. Hollow: believed; solid: measured (95% CI in 2025; follow-up level not stated).

Measure what students can do after AI help.


RECENT RESEARCH • TWO LARGE REAL-WORLD STUDIES AND TWO RANDOMIZED COLLEGE TRIALS

1

STRÖMBERG ET AL. • 26,811 SECONDARY STUDENTS • 30 MONTHS

[1]

Change after students began using AI (% of the pre-AI average)



Losses were concentrated among the ~80% of AI users whose homework time collapsed; those who kept spending the time scored about the same.


*Students using AI 2+ years (avg. -7%). Not randomized; one county in central China, 1M+ people.

2

RISMANCHIAN ET AL. • ALEKS MATH • COLLEGE STUDENTS

[2]

After ChatGPT, on word problems that can be pasted into a chatbot: compared with interactive graph problems that can't be pasted (the control)



Less practice, less retained knowledge.

The graph problems show what would have happened without easy AI help, so the gap is the change linked to AI.

Bars show 95% CIs. Not randomized; under review; McGraw Hill co-authors.

3

Randomized college trials: design matters

[3]

Maryland, 2,379 undergraduates: adding a course AI tutor cut participation (-0.90 SD); same-course grades -0.37 SD.

Online MBA, 86 students: a tutor built to make students reason improved written reasoning (+6.6 of 55 points; next slide).

[4]

When AI does the work, practice looks better and knowledge drops.

Students who kept doing the work learned about as much. Test yourself later, unaided, on a changed problem.

Sources: [1] Strömberg et al. (2026) [2] Rismanchian et al. (2026) [3] Liu, Sweet et al. (2026) [4] Yang et al. (2026)


Left: point estimates (intervals in the paper). Right: 95% CIs, each relative to the graph-problem control. Different designs; not pooled; more studies in the notes.

Attempt first, then ask AI for hints, not answers.

TWO RECENT RANDOMIZED TRIALS IN REAL COURSES

1 ECONOMETRICS COURSE • FALL 2025 • 289 STUDENTS [1]


Everyone had the course’s ChatGPT assistant.
Half also got three emails with tutor-mode prompts.



In-person exam; the gain came on open-response problems.

2 ONLINE MBA • SUMMER 2026 [2]

A GPT-5.4 tutor required an attempt and gave a smaller step, never the answer. 86 students.



Written answers that build a connected argument; gap +22 points (95% CI 2 to 42).
51 of 86 took both tests; AI graders scored answers.

3 THE PROMPT THEY GOT

“Act as an econometrics tutor for my exam. Do not give me the final answer. ... Guide me step by step, asking me to complete the next step myself. At the end, give me a similar exercise to solve on my own and correct it.”

Other emails: practice questions without solutions; then check the AI’s own answer.

Add your topic, notation, and your own attempt first.

[1] Abridged from the trial’s first email.

The first obstacle is a reason to think further – not an automatic hint.
Seeking help? Give your course topic and notation, your attempts, the precise gap, and the help allowed.

Sources: [1] Gallegos (2026) [2] Yang et al. (2026)

Pass bars: control rate and control rate + estimated effect [1]. Both are preregistered working papers; not every study agrees (see notes). Follow course rules.

Check the rules before using AI or sharing data.

1 Ask in your course [1]

2 Follow your employer’s policy

USES	Which uses does this assignment permit?	→ Use only approved AI tools, for approved tasks.
UPLOADS	May I upload these course materials?	→ Share data only through approved services.
OWN WORK	What work must I do independently?	→ Obtain the required review before release.
DISCLOSURE	How should I disclose AI assistance?	→ Document AI assistance in the project record.

3 IF ASKED TO CONCLUDE MORE THAN THE EVIDENCE SUPPORTS

State the limitation and specify the evidence required.

Consult your supervisor when necessary. Correct errors promptly, including after delivery. [2]

Following course rules builds the habit employers require.

If permission is unclear, resolve it before acting. Access to a tool is not permission to use it.


Sources: [1] [Purdue teaching guidance ↗](#) [2] [ASA ethical guidance ↗](#)

Course and employer requirements vary; the workplace column is general guidance, not any specific employer’s policy.

Build the expertise to check and steer AI, step by step.

Why protect your learning?

SUGGESTED PROGRESSION • NOT A PURDUE RULE • COURSE RULES APPLY AT EVERY STEP



New topic at any stage? Start again at the first step.

Expect employers to keep the people who can check and steer AI.
Move up a step when you can explain the idea and solve a changed problem without AI — not when you feel confident.

Sources: [\[1\] Gallup \(2026\)](#) [\[2\] Brynjolfsson et al. \(2026\)](#) [\[3\] Stanford Canaries tracker](#) [\[4\] Dratsch et al. \(2023\)](#) [\[5\] Sinha & Kapur \(2021\)](#) [\[6\] Palazzo et al. \(2010\)](#) [\[7\] Eloundou et al. \(2024\)](#)

AI-exposed: the two most exposed fifths of occupations by GPT-4 task ratings [7], tracked monthly [3]. Job trends are descriptive; causes debated. Suggested progression, not a Purdue rule.

Small course topics can matter long after the exam.

FOUR DOCUMENTED PATHS • COURSE TOPIC → TOPIC BEYOND THE COURSE → AI

IN YOUR COURSES [1]

1

MA 16100 / 16500 • Calculus I

The chain rule

2

STAT 41600 • Probability

Conditional probability

3

MA 35100 • Linear algebra

Vectors and dot products

4

STAT 35500 • Statistics for data science

Confidence intervals

BEYOND THE COURSE

Automatic differentiation

Software applies the chain rule to every step of a program; backpropagation is a special case. [2]

Markov chains

The next state depends only on the current one; later course: MA 43200. [4]

Cosine similarity

$\cos \theta = \frac{u \cdot v}{\|u\| \|v\|}$: do two vectors point the same way?

Conformal prediction

Sets with guaranteed coverage, such as 90%, for any model's predictions. [8]

IN AI

Training a neural network

Backpropagation works backward through the network by the chain rule; training updates the weights with its gradients. [3]

Predicting the next word

N-gram models: $P(\text{next word} \mid \text{last few words})$ [5]
GPT: $P(\text{next token} \mid \text{previous } k \text{ tokens})$, k up to 512 [6]

Finding relevant passages

Text embeddings are compared by cosine similarity to find the closest passages. [7]

Uncertainty for AI answers

Sampled answers are kept until the set covers an acceptable one, with chosen probability. [9]

Keep what you learn; it may matter in ways no syllabus mentions.
The path from a course to AI runs through topics you study next. Ask what each topic leads to.

Sources: [1] [Purdue DS–Statistics plan \(2025–26\)](#) ↗ [2] [Baydin et al. \(2018\)](#) ↗ [3] [Rumelhart et al., Nature \(1986\)](#) ↗ [4] [Purdue MA course descriptions](#) ↗ [5] [Jurafsky & Martin, ch. 3](#) ↗ [6] [Radford et al., GPT \(2018\)](#) ↗ [7] [Reimers & Gurevych \(2019\)](#) ↗ [8] [Angelopoulos & Bates \(2022\)](#) ↗ [9] [Quach et al. \(2024\)](#) ↗

Course codes from the 2025–26 guides [1, 4]; each course is required or an elective in Purdue Statistics or Data Science majors (notes list which). Examples, not a full account of AI.

Use AI to connect your courses to your goals.

SUGGESTED ROUTINE • INDUSTRY EXAMPLES AND EVIDENCE FROM PUBLISHED SOURCES

1 Name the topic and your interest

“STAT 41600 just covered conditional probability and Bayes’ theorem. I want to work in sports analytics. Where does this topic matter there? Give two concrete examples and the later Purdue courses that build on it.” [1]

USED IN INDUSTRY

MLB publishes xBA: $P(\text{hit} \mid \text{exit velocity, launch angle})$ [2]
Opta xG: $P(\text{goal} \mid \text{shot situation})$, adopted by clubs and broadcasters [3]

2 Check what it tells you

Confirm courses in the catalog; open every source it cites.

3 Write the connection yourself


A few sentences in your own words: why it matters to you.

4 Follow one thread

A dataset, a reading, an elective, or a question for office hours.

Why write it yourself?

88 undergraduates learned a mental-math method. Less-confident students: problems solved correctly vs. students given no reason.



Laboratory study; effects are Cohen’s d for students one SD below average confidence. [4]

For you: when a topic isn’t clicking yet, write your own reason; an AI’s explanation is being told.

No course can cover every interest; connecting the dots is up to you.

Use AI to find connections, then check them, write them yourself, and follow one.

Sources: [1] [Purdue STAT course descriptions](#) ↗ [2] [MLB Statcast glossary \(xBA\)](#) ↗ [3] [Opta Analyst \(2023\), xG](#) ↗ [4] [Canning & Harackiewicz \(2015\)](#) ↗

Suggested routine; [2, 3] show the sample topic in industry use. The evidence concerns writing about relevance, not AI specifically. Follow course rules.

Know the kinds of AI tools and when each is safe to use.

KINDS OF AI TOOLS • FROM VENDOR TERMS AND DOCUMENTATION

		What it can do →			When it is safe to use
		Answer in a chat	Run from your code (API)	Act in your project (agent)	
Whose terms it runs under →	Personal accounts free or personal plans	ChatGPT, Claude, Gemini personal sign-in	Free API keys [1] e.g., the Gemini API free tier	Coding agents [2] [3] [4] Claude Code · Codex · Gemini CLI	Public data only; what you send may be kept, reviewed, or used for training. [11] [12]
	Business terms work, school, or paid API	Work versions [5] e.g., Copilot with a work account	Paid API keys [1] e.g., the Gemini API paid tier	The same agents [2] [3] [4] on a work plan or paid API key	Data your organization approves; business terms typically exclude training. [3] [5]
	In-house servers e.g., Purdue GenAI Studio	GenAI Studio chat [6] Log in: genai.rcac.purdue.edu ↗	GenAI Studio API [7] the same models from Python	Build your own agent [8] with the GenAI Studio SDK	What the organization allows; at Purdue, sensitive but not FERPA. [6]
	Your computer local, open weights	llama.cpp · Ollama [9] [10] chat with a local model	Local server [9] scripts call a local model	—	Data stays local unless a cloud feature sends it out. [10]

The same tool can sit in different rows: the account or plan you use sets the terms.

Know whose terms a tool runs under before you share data.

Personal accounts: public data only. Business terms and in-house servers: what your course or employer approves.


Knowledge bases match meaning; databases match exactly.

REAL CODEBOOK PASSAGES AND RELEASED RESULTS • SIMILARITY COMPUTED FOR THIS SLIDE

1 Knowledge base • ranks passages by meaning

Each passage and your question → embedding → rank by similarity [1]

Q: How did they measure players' skin tone?



Different words still match: “measure” found “coded by two independent raters.”
You check: does the closest passage answer your question?


✗ Wrong tool: ask the knowledge base “How many red cards were given in total?” Its top match (0.80) is the redCards definition, not a count. [3]

Search documents for meaning; query records for exact answers.
You choose the tool, read what comes back, and decide what each condition should mean.

2 Database • checks every row

A table of 4,194 runs → SQL conditions → every row tested [4]

Q: How many acceptable TV-news analyses said “supported”?



Let coding agents act only within limits you set.

SUMMARIZED FROM VENDOR DOCUMENTATION (OCTOBER 12, 2026) • NOT A SECURITY TEST

1 The agent proposes	2 The limits decide	3 What happens
Read files in your project	→ read-only work inside the sandbox [1]	→ RUNS without asking
Edit code or run a command	→ Claude Code, manual mode: ask first [2]	→ ASKS YOU approve or decline
Reach the internet	→ Codex: network off by default [1]	→ BLOCKED unless you approve

Then review: read the changes, run the tests, keep a checkpoint.

Prompt injection: a file or web page can hide instructions that try to steer the agent; the limits above still decide. [4]

4 Learning cost • randomized study of 52 programmers [3]	The AI group scored lower on the quiz; the patterns below were observed, not randomized.
Relied on AI to write or debug code: quiz average below 40%	Asked about concepts or for explanations: 65% or higher

You do the thinking; let the agent type only what you understand.
Keep approvals on, review every change, and follow your course’s rules.

Try one new way to use AI this semester.

Pick one route and try it on something small, with public or synthetic data.



COMPARE, THEN SCRIPT IT

Purdue GenAI Studio

Send one prompt to two models side by side. Then send it from a short Python script through the API.

Log in: genai.rcac.purdue.edu ↗ API guide ↗

Approved for public and sensitive (L2) data, not restricted data such as FERPA or HIPAA records. [1] Keep your API key out of your code.



SEE WHERE IT RUNS

A model on your computer

Run a small open-weight model with llama.cpp or Ollama; either can serve it to your scripts.

[llama.cpp](#) ↗ [Ollama](#) ↗

Choose a model your hardware can run, and check whether your setup sends data anywhere else.



BUILD, THEN REVIEW

A coding agent

Try Claude Code or Codex on your own project, with approvals on. Read every change and run the tests.

[How agent approvals work](#) ↗

For coursework, follow the course’s rules. Test any tool like a statistician: checkable questions, repeated runs, recorded settings.

WHERE TO GO NEXT Now: CS/SCI 173, the College of Science AI Competency course [5] Later: Chapter 6, LLMs in data science [6]


Learn one new way to use AI this semester, and keep the thinking yours.
Record the model, date, settings, and your checks so someone else can repeat the work.

Sources: [1] [Purdue RCAC, GenAI Studio \(Sept. 29\)](#) ↗ [2] [Purdue RCAC, API guide](#) ↗ [3] [llama.cpp](#) ↗ [4] [Codex security docs](#) ↗ [5] [Purdue Science, CS/SCI 173](#) ↗ [6] [Computational Data Science, Ch. 6](#) ↗

RCAC documentation updated September 29, 2026; check current access, limits, and data rules before you start.

Build what interests you with AI, then learn how it works.

SPEAKER'S SUGGESTION • YOUR CHOICE OF PROJECT • EXAMPLE FROM MLB SOURCES



EXAMPLE • IF YOU LIKE BASEBALL

MLB's expected batting average: $P(\text{hit} \mid \text{exit velocity, launch angle})$, from comparable batted balls [1]
MLB's public Statcast data list both measurements, each play's result, and MLB's own estimate [2]
Build a model with AI → learn what it uses → compare yours with MLB's.

Let a project you care about show you what to learn next.

Personal projects can use a lot of AI help; coursework follows each course's rules, and the learning stays yours.

Sources: [1] MLB Statcast glossary (xBA) ↗ [2] Baseball Savant CSV fields ↗

Optional strategy, not a validated intervention. Example definitions: MLB Statcast glossary [1] and Baseball Savant CSV fields [2]; no model was built for this talk.

You do the thinking; let AI support your growth at Purdue.

MAIN POINTS • THE FOUR QUESTIONS FROM THE START

Purpose

What is the purpose?

- Match AI use to what you need to learn.
- Write your own connection from courses to goals.

Permission

What is allowed?

- Follow course rules now, employer policy later.
- Share data only where a tool's terms allow.

Reasoning

What must I do myself?

- The thinking: attempt first, then ask for hints.
- Build the expertise to check and steer AI.

Checking

How will I check it?

- The same data can lead AI to different conclusions.
- Read the passage, check the query, review every change.

THE THINKING IS YOURS • SPEAKER'S ADVICE

Give a hard proof or project days, not minutes.

Struggling with it is how you develop. Don't be quick to turn to AI for the quick fix.

Across 53 studies, trying problems before being taught built deeper understanding. [1] AI users who kept spending the time scored about as well as students without AI. [2]

START THIS WEEK

Pick one hard problem from your coursework and check what help is allowed. Stay with it over several days, then bring one precise question to an instructor or TA.

Questions?