

AI NEXT WAVE · FREE DOWNLOAD

The AI Check Kit

Checklists for AI Headlines, Hiring, Money, and Privacy

Separating AI fact from AI fiction.

By Jay Cadmus, AI Next Wave

AI is screening your resume, scoring your loan, listening at the doctor's office, and watching your porch. This kit gives you the exact questions to ask, the settings to check, and the steps to take, in plain language. Every checklist first ran in AI Next Wave, where I check AI headlines against the sources.

What's in the kit

1. Use AI Without Getting Played: 11 Prompt Tricks That Actually Work
2. When Software Screens Your Resume: Five Moves Before You Apply
3. Turned Down for a Loan? What to Ask About the AI Score
4. If a Tool Holding Your AI Keys Gets Breached: Four Steps
5. Four Checks If You Use an AI Browser
6. Your Doorbell Camera: Six Settings Checks
7. At the Doctor's Office: Three Questions About AI Recording
8. Homework in the AI Era: Three Checks for Your Household
9. When a Data Center Comes to Town: Three Things to Watch
10. Bonus: The Before Picture Worksheet

Work through them in any order. Each one stands on its own and takes a few minutes.

Use AI Without Getting Played: 11 Prompt Tricks That Actually Work

I use AI tools every day for analyst work. Most prompt advice is fluff. These 11 are the ones that survived. Each one comes with the exact words to type.

1. Draw a line between context and instructions

Models blend everything together. A --- line tells it where background ends and the job begins.

```
[ your pasted context here ]  
---  
Summarize the above in 5 bullets for a busy manager.
```

2. Show a bad example too

"Do this, not that" calibrates the model twice as fast as praise alone.

```
Bad: "Revenue was good."  
Good: "Revenue was $2.1M, up 4% vs last quarter."  
Write like the good one.
```

3. Make it plan before it acts

A bad plan executed well is still a bad result. Check the plan first.

```
Outline your approach in 4 steps.  
Wait for my go-ahead before doing anything.
```

4. Use a magic word

Loading files or pasting chunks? The model jumps the gun. Tell it to hold until you say so.

```
Do not process or summarize anything until I say the word PROCEED.
```

5. Chain prompts, check between steps

One mega-prompt means one mega-failure. Review each step before the next one runs.

```
Step 1: pull the key numbers from this sheet.  
I will review them. Then we move to step 2.
```

6. Give it a rubric, not vibes

"Make it better" gets random changes. Scored criteria get targeted ones.

Score this draft 1-5 on clarity, accuracy, brevity.
Then fix whatever scores lowest.

7. Ask what you're missing

The highest-leverage question. It surfaces angles you didn't think to ask about.

What am I missing?
What should I be asking that I'm not?

8. Force the shape of the answer

Structure is thinking. Fix the shape and the logic has to fill it.

Answer in exactly three sections: what happened, why it matters, what to do next.

9. Demand confidence per claim

Shows you exactly where to verify instead of re-checking everything.

Rate your confidence 1-10 on each claim in your answer.

10. Let it interview you

For fuzzy tasks, answering questions beats writing a perfect brief.

Ask me one question at a time until you can scope this task. Then summarize your plan.

11. Lock it to your source

The single best hallucination killer. No source, no answer.

Use only the notes pasted below.
If the answer is not in there, say so.

When Software Screens Your Resume: Five Moves Before You Apply

A computer may reject your resume before a person sees it. Derek Mobley applied to more than 100 jobs through companies using Workday's hiring software. He got zero interviews. Workday's own court filings say its tools rejected 1.1 billion applications during the period in question.

A Stanford study from June 2026 found the same pattern across the industry. The honest complication: tested AI can judge more evenly than a hiring manager's gut. The difference is testing.

1. Assume software reads your resume first. Match your resume's words to the actual skills in the posting; applicant-tracking filters are the first gate.
2. Ask the employer directly: "Does AI screen applications for this role, and what does it evaluate?" In New York City, companies owe you at least 10 business days' notice that an automated tool will be used, and you can ask for an alternative process. In Illinois, employers must tell you when AI is in the mix, and the law gives you a way to request human review. In California, a 2026 law (AB 331) lets workers request which factors the AI considered.
3. If a rejection smells off, ask for a human second look. The NYC alternative-process right and the Illinois human-review requirement exist for exactly this.
4. Keep records. Screenshot the posting, the application confirmation, and any AI disclosure. If bias claims need proof, documentation wins.
5. Don't rely on tricks. About 1% of real resumes in a 2026 USENIX Security study carried hidden prompt injections aimed at AI screeners, and vendors are now building detectors for them. A tailored, honest resume beats a hacked one.

You still have to play the game. You have the right to know the rules.

Turned Down for a Loan? What to Ask About the AI Score

A lender, Earnest Operations, paid \$2.5 million to Massachusetts in July 2025 to settle a lawsuit. Massachusetts said the company's AI used a score built partly on your college: how often students there paid back their loans. The state also said Black and Hispanic applicants were more likely to be turned down, or offered worse deals, because of that score. Earnest denied the claims and denied breaking the law, and a settlement is not a ruling.

What protects you, and what to do:

1. Know the rule. In May 2026, the CFPB, the government agency that watches banks and lenders, said it in writing: even when a computer makes the decision, the lender has to give you the specific, real reason you were turned down, even when its model is hard to explain.
2. If you get turned down, ask for the exact reasons in writing.
3. Read that letter closely. That letter is how you spot bad data and fix it.

If a Tool Holding Your AI Keys Gets Breached: Four Steps

An API key is a private code. A tool holding it can run AI and bill the use to the account owner. OpenCode said on Oct 17, 2026 that someone got unauthorized access to an older system, and keys users had saved in its Console workspaces may have been exposed. It told users to cancel keys on the provider site and watch for unexpected usage.

If you get an email like that, do this, on the provider site itself:

1. Skip the links in the email. Type the provider's address yourself.
2. Cancel the old keys and make new ones.
3. Check usage back to the date the company names (Oct. 17 in the OpenCode case) for charges you don't recognize.
4. Set a spending limit, a cap that stops charges past an amount you pick.

Four Checks If You Use an AI Browser

Your AI assistant may know more about your browser history than your family does. Run these four checks:

1. Turn browser memory off in settings.
2. Never sign into banking, email, or work accounts inside an AI browser.
3. Remove extensions you do not recognize. Treat "read and change all your data" as a red flag.
4. Assume anything you type in the sidebar is collected.

Your Doorbell Camera: Six Settings Checks

Your camera does two jobs. Event spotting flags a person, a pet, a car, or a package, with no name attached. Face naming measures a face, turns it into a set of numbers called a faceprint, and matches it against saved faces. You can keep the first and skip the second.

1. Leave face naming off. Keep the person alert on.
2. Open the saved faces list now and then. Delete faces you no longer need.
3. Shorten how long clips are kept, if your app offers the choice.
4. Check the sharing page. Know who can request your video.
5. Aim the camera at your porch. Keep the public sidewalk out of the frame where you can.
6. Turn on two-step sign-in for the camera account.

At the Doctor's Office: Three Questions About AI Recording

An ambient scribe records the visit through a phone or mic in the exam room, and the AI writes the clinical note from the transcript. The audio goes to the vendor's servers, not just your doctor's system, and vendors set their own retention rules. Ask before the exam starts:

1. Is a scribe recording this visit, and where does the audio go?
2. How long is it kept, and can I say no?
3. Can you turn it off if I want to say something private?

One more step: read your note in the patient portal and ask for corrections. A real-world study of 356 AI notes found 18% had omissions and 11.5% had hallucinations. The doctor is legally responsible for every signed word, and a wrong note can follow you.

Homework in the AI Era: Three Checks for Your Household

In a study of 26,811 students, homework took 30% less time and scored 18% higher after AI adoption, while closed-book exam scores fell 20% within six months. The shortcut is the problem, not the tool; students who kept spending the same time on homework learned just as much as everyone else.

1. Homework scores are a broken signal now. High scores can mean the AI did the work.
2. Watch homework time. If the same assignment takes half as long with a better grade, the learning probably did not happen.
3. Ask the school what its AI policy is. Most still do not have one.

When a Data Center Comes to Town: Three Things to Watch

1. Where a data center sits matters more than the national total. About 40 percent of US data centers are in high or extreme water-stress areas.
2. If you live somewhere dry and a center is going up nearby, watch the water permits. That is where the household cost shows up first.
3. Save this for the next town hall.

The Before Picture Worksheet

For your work or your business. A good workflow audit needs a before picture. Before I recommend a fix for a slow process, I capture the work as it runs today. Five fields, every time. The fields are blank by design: fill them in for one process that feels slow, stuck, or easy to drop.

1. The pain.

Write it in one sentence a stranger could follow. Name the job, the delay, and who waits at the end of it. Start where the work feels slow, stuck, or easy to drop.

2. The current path.

List every step in the order it really happens, from the trigger that starts the job to the hand that finishes it. Mark each handoff. Include the typing people do twice: the same names, dates, or numbers entered in two places.

3. Elapsed time.

How long the whole job takes from start to finish, waits included. That total is the clock the person waiting actually feels. Count the hands involved too; every extra person adds a handoff, and every handoff adds waiting.

4. Error repair.

The mistakes someone has to catch, fix, and send back before the job can move again.

5. Cost of waiting.

What sits on hold while this job sits still: the next job, the answer a customer is owed, the decision on hold.

Then run the smallest test first. Two rules hold the whole thing together. Call every early number a guess until you have timed the work yourself. And set the success threshold before the test starts: the number that has to move, and by how much, for the fix to stay.

KEEP GOING

Get AI Next Wave free by email

This kit is one morning's worth of checking. AI Next Wave is the daily version: one short edition, every morning, checking the day's AI headlines against the sources and covering what they mean for your job, your money, your safety, and your privacy.

Subscribe free

theainextwave.com/subscribe · One short edition every morning. No hype, no spam, unsubscribe any time.

Watch and listen

YouTube: youtube.com/@jaycadmus, daily AI fact-check Shorts and a weekly long-form video.

Podcast: AI Next Wave on Spotify, the week's AI stories in about six minutes, every Monday, Wednesday and Friday.

Instagram and TikTok: [@thejaycadmus](https://www.instagram.com/thejaycadmus). Threads: [@thejaycadmus](https://www.threads.net/@thejaycadmus).

Share this kit

The kit is free. If it helped, send a friend to theainextwave.com/check-kit/ so they can get their own copy and join the list.

I check AI headlines against the sources.

Jay Cadmus, AI Next Wave

Separating AI fact from AI fiction.