CAN YOU DEFEND EVERY WORD?

A Constraint Design Guide and Starter Path for Career Educators

Designing work that cannot be outsourced, and tools that hand it back



Megan O'Connor, SHRM-CP, CPRW

Talent Transformation Whisperer, LLC | transformwithttw.com

Audience: Career Services Professionals | Adaptable for student workshops, course modules, and advising tools

ABOUT THIS GUIDE

I wrote this guide for people who design career readiness instruction: a credit course, a workshop series, a required module, or the advising curriculum that does the same work without a course number. It teaches a method, not a product. It does not reproduce the instruction set behind any particular coaching agent, and every example below was written for this guide as a pattern to adapt rather than lifted from a live system.

The problem it addresses is specific to our field. Career readiness runs on a small set of deliverables, and generative AI is better at producing all of them than it is at almost anything else. A resume, a cover letter, a LinkedIn summary, a brand statement, and an interview answer are precisely the artifacts these tools were trained to generate. If we keep assigning them the way we always have, we are no longer teaching career readiness. We are collecting output.

Nine parts across three tracks: the course, the assignment, and the agent. Pick your track below and read only what you need.

THREE TRACKS

There are three places to put a constraint, and most people arrive needing only one of them. Use whichever track matches what you are actually building this term.

Building a course or a workshop series?

Start at Part 1, then use Part 2 to pressure test the deliverables you already assign.

Revising one assignment?

Go straight to Part 2, find your deliverable, then pick a constraint from Part 3.

Building an AI coaching agent?

Parts 4 and 5. Read Part 1 first anyway, because an agent built before the assignments are sound will be asked to compensate for a design problem it cannot fix.

If you picked this up at a session:

The four sections we walked through on screen are the hard boundary, refusal and redirect together, the repeat ask, and the closing check. Part 4 has the other nine. Part 6 is the break test, which takes twenty minutes and is the fastest way to find out whether your own agent holds.

THE IDEA IN ONE PAGE

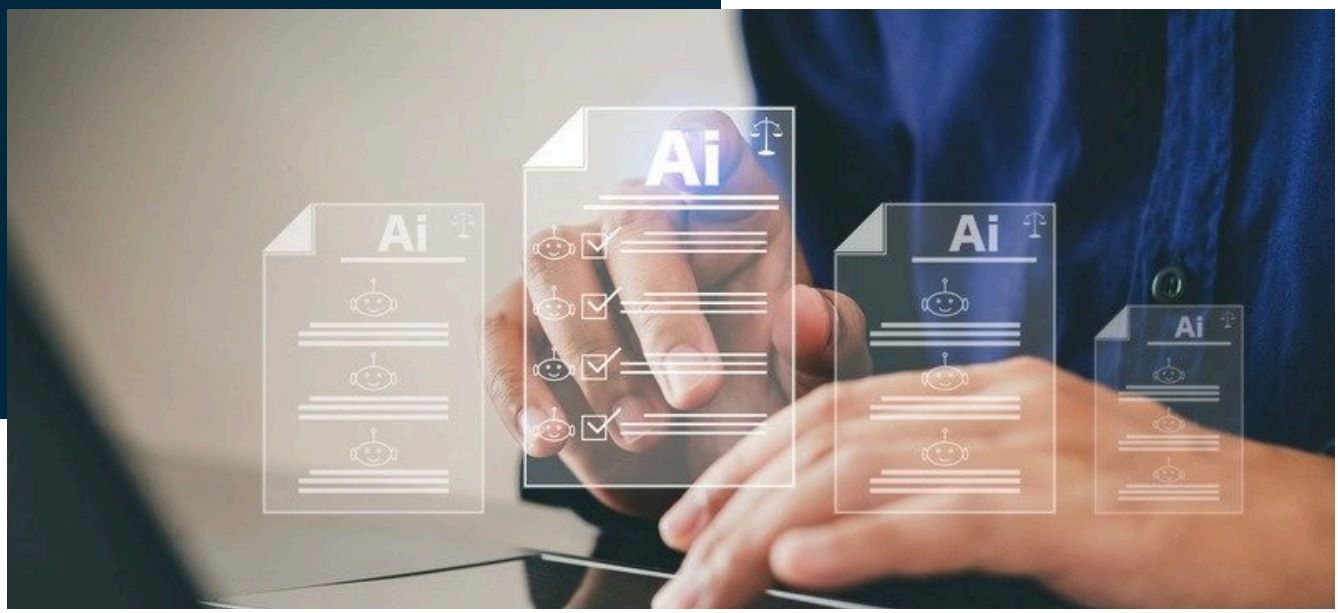
Most campuses respond to AI with a policy. A policy tells students what they may not do, and then asks them to want to comply. It is worth asking what that policy is protecting if the assignment can be completed by pasting it into a chatbot.

A constraint is different. A constraint makes the shortcut unavailable, so compliance stops depending on the student's willpower or their memory of a rule. Constraints live in two places, and the strongest designs use both. An **assignment constraint** changes what the work requires. A **tool constraint** changes what the AI will do when asked. Neither one alone is enough: a well designed assignment can still be undermined by a tool eager to complete it, and a carefully instructed tool cannot rescue an assignment that was always outsourceable.

Every constraint in this guide serves one governing question, could you defend every word of this in an interview?, and the GPS Principle behind it: AI navigates, the student drives. A constraint is how that principle becomes part of the work itself.

THE QUESTION THAT STARTS EVERYTHING

What does this work require that only this student can supply? Write that down first. The constraint is just the sentence that makes the answer unavoidable.



Part 1.

DESIGNING THE COURSE

Six decisions that shape a career readiness course before you write a single assignment. They apply equally to a credit course, a workshop series, or an advising curriculum that does the same work without a course number.

1. Put the commitments in week one, before anything is produced

The first week of a career readiness course is usually orientation and a syllabus quiz. Use it. Teach what AI may and may not do in this course, what a hallucination costs when it invents an employer, and the question that governs everything: can you defend every word of this in an interview? Nothing is at stake in week one, which is exactly why the ethical foundation holds there. By week four the student has a deliverable due and the conversation is harder.

2. Assign at least one artifact that requires a real human being

Every career readiness curriculum already has candidates: an informational interview, a networking event, a mentor conversation, a mock interview. Make one of them required and make the write up depend on what the other person actually said. This is the single most AI resistant assignment available to our field, and most of us already have it in the syllabus.

3. Let one artifact accumulate across the term

Career readiness deliverables are usually independent, which makes each one individually outsourceable. Pick one thread and build it weekly: a presentation assembled a slide at a time, a single job posting that returns for applicant tracking analysis and then for tailoring, a portfolio that grows. Accumulated work cannot be produced the night before, and it shows you drift early enough to coach rather than catch.

4. Escalate the demand, do not repeat it

If your institution offers career readiness instruction at more than one level, the second encounter should ask something harder than the first, not cover the same ground with the same artifacts. One useful progression: a first course asks whether a student can explain their work out loud, and a later course asks whether they can defend every word of it in an interview with their own name attached.

5. Decide what AI may touch in each unit, and write it into the instructions

A course level AI policy is too blunt. The right level of AI involvement differs by unit: brainstorming a brand statement is a reasonable use, drafting a thank you note for an interview the tool did not attend is not. Make the call unit by unit, say it in the assignment itself, and say why. Students follow instructions that come with a reason far more reliably than rules that arrive as a list.

6. Build assignment constraints before you build a tool

A coaching agent built before the assignments are sound will be asked to compensate for a design problem it cannot fix. Fix the assignments first. They will also tell you what the tool actually needs to refuse, which is the only reliable specification you will get.

Part 2.

THE CAREER READINESS ARTIFACT AUDIT

Every career readiness curriculum runs on roughly the same set of deliverables, and generative AI produces all of them well. This table names where each one is vulnerable, the constraint that closes the gap, and what the student is then the only possible source for. Start with the deliverable you assign most often.

DELIVERABLE	WHERE AI TAKES OVER	THE CONSTRAINT THAT CLOSES IT	WHAT ONLY THEY CAN SUPPLY
Resume	Generated wholesale from a job posting, with invented or inflated bullets the student cannot explain.	Require the student to run their own resume against a posting they selected, then write about the match rate, what was missing, and what they will revise.	Their actual experience, and a judgment about what to include.
Cover letter	Produced in seconds, polished, interchangeable, and unconnected to any specific conversation or company.	Require a draft before AI, then AI for review only, with a screenshot of the prompt submitted alongside.	The reasoning for why this employer and this role.
LinkedIn profile	Written entirely by AI in a voice that does not match the person a recruiter later meets.	Require the live profile URL plus a short account of what the student changed from any AI suggestion and why.	A public presence that matches the person behind it.
Job or internship search	A list of postings generated without the student ever using a job board or reading a description closely.	Require a real posting that demands their actual major and the experience they will genuinely hold at graduation, captured as text.	Honest self assessment against a real target.
Career exploration	A plausible summary of an occupation, assembled from training data, with no verification and no personal fit.	Require verification against the original source, plus questions only the student can answer: commute, cost of living, daily rhythm, what their life looks like outside the work.	Whether this life is one they want.
Networking outreach	A message drafted and sent verbatim to a stranger, in language the student would never use in person.	Require the student's own draft first, with AI limited to checking tone, and require them to report the response they received.	A relationship, which no tool can build.
Informational interview	A set of generic questions, and a write up that could have been produced without the meeting happening.	Require the name of the person, how the student found them, what was discussed, and what the next step in the relationship will be.	A conversation that actually occurred.
Interview preparation	Scripted answers memorized rather than developed, which collapse at the first behavioral follow up.	Require a live unscripted mock interview. AI may help prepare; it cannot sit in the chair.	Performance in real time, in front of a person.
Thank you note	Written by a tool that was not in the room, referencing nothing that happened in it.	Require it within 24 hours of a real interview, referencing one specific moment from that conversation.	A detail only someone present could know.

DELIVERABLE	WHERE AI TAKES OVER	THE CONSTRAINT THAT CLOSES IT	WHAT ONLY THEY CAN SUPPLY
Personal brand statement	Fluent, impressive, and unrecognizable as the student when read aloud by someone who knows them.	Permit AI explicitly, and say plainly that it will need information about the student to produce anything that sounds like them. Then require revision until it does.	Their own voice, tested against someone who knows it.
Reflection	The easiest deliverable in the field to outsource, and the one most often assigned without constraint.	Ask about experience, reaction, and judgment: what surprised you, how did you feel, what will you do differently and why.	An interior life the tool does not have access to.

Reading the last column:

If you cannot fill it in for one of your own assignments, the problem is the assignment rather than the tool. That column is the assignment's reason for existing.

Part 3.

SIX CONSTRAINT TYPES

The audit above applies these six. You will rarely need more than two in one assignment. Two well chosen constraints do more than six stacked together, and a student facing six reads them as suspicion rather than design.

Irreducible source

The work must reference something only this student witnessed or did. A tool cannot fabricate a conversation that happened in a specific room on a specific day.

PATTERN

Write a follow up message to the person who interviewed you, referencing one specific moment from your conversation.

Watch for the generic escape hatch. If the student can satisfy the requirement with a plausible invention, tighten it until they cannot.

Sequence

The student's own thinking must exist before AI touches the work. This preserves the drafting effort that builds judgment, while still allowing the tool to help.

PATTERN

Write your first draft before using AI. You may then use AI to review tone, grammar, clarity, and structure. AI may not write the draft.

This is the single most portable constraint in the guide. It fits almost any writing assignment and requires no new technology.

Documentation

The student records their own AI use. This makes the process visible without surveillance software and without the instructor playing detective.

PATTERN

Submit a screenshot of the prompt you used, alongside your final work.

Frame it as professional practice rather than proof of innocence. Documenting your sources is a credibility habit, not a confession.

Reality match

The artifact must match verifiable facts about this student. The constraint forces honest self assessment before anything is drafted.

PATTERN

Select a posting that requires a degree in your actual major, and experience you will genuinely hold at your graduation date.

This one surfaces mismatches early, while there is still time to coach, rather than at the point of rejection.

Accumulation

The work builds visibly across the term, so the record of the student's thinking cannot be generated the night before.

PATTERN

Build one slide of your final presentation each week, beginning in week two.

Accumulation also makes drift visible to you. A week that suddenly does not sound like the previous six is a coaching conversation, not a violation.

Questions only they can answer

Affective and metacognitive prompts ask about experience, reaction, and judgment. A tool can guess at these, and the guess reads as hollow to any reader who knows the student.

PATTERN

What surprised you? How did you feel during it? What will you do differently, and why?

These questions also carry the part of learning that cognitive frameworks tend to miss, which matters more in AI assisted work, not less.

Part 4.

BUILDING THE AGENT

An agent is only as good as the behavior you can describe in advance. The sections below are what a working instruction set needs. You are not writing a prompt. You are writing a policy that a tool executes consistently, at two in the morning, when you are not there.

Write it in plain language, in your own voice. These tools follow prose instructions better than they follow bullet lists of rules, and your students will hear whatever tone you wrote.

The section most people leave out is the repeat ask.

Without it, an agent gets firmer each time a student asks again, and the student who needed help most is the one who gets the coldest response.

SECTION	WHAT IT NEEDS TO CONTAIN
Identity and purpose	Who the agent is, which course it serves, and the two or three capacities it exists to build. Name them explicitly, for example, competencies from the NACE career readiness framework (National Association of Colleges and Employers, n.d.). An agent told it builds judgment behaves differently from one told it helps with assignments.
Who your students are	A plain statement of the population and the stance you want taken toward them. Treat every student as capable. No deficit framing, no urgency pressure, no shaming language. This shapes tone in every exchange that follows.
The governing principle	One sentence the whole agent answers to, plus the single question that tests it. Everything else in the instructions is downstream of this.
The hard boundary	A named list of what the agent will never produce. Write it by type of writing, not by assignment title, so a student who renames the request has not found a loophole. State that it does not bend to pressure, flattery, urgency, or claims of permission.
Refusal and redirect together	Instruct the agent to decline and coach in the same message, never a bare no. Tell it to name what the student is actually trying to accomplish, then ask for the one specific thing only they can supply.
The repeat ask	What to do when a student asks twice. A student asking again is usually stuck, short on time, or worried they are doing it wrong. Treat that as a signal to help, not a rule to enforce harder. Hold the boundary and move to the real obstacle.
Name the move	Tell the agent to explain why it declined, in coaching terms rather than policy terms. Students who hear the reason internalize a standard. Students who hear a rule learn to route around it.
Who moves first	The agent opens by asking what the student has drafted, never by producing content. If they have nothing, it coaches them to a first sentence rather than offering one.
How it coaches	Questions before answers. React to what the student produced rather than replacing it. Point to the spot that needs revision without supplying the words. Push generic language toward the specific detail underneath.
The evaluation lenses	The questions the agent asks when a student brings it AI output. Instruct it to ask the questions and let the student answer, rather than evaluating on their behalf.

SECTION	WHAT IT NEEDS TO CONTAIN
Voice and referral limits	Describe the tone you want in plain words, including what to avoid. Then name what the agent is not: not a counselor, not an authority on your policies. Tell it where to send a student in distress.
Context and calibration	A week by week map of what students are working on, so the agent recognizes where they are. Plus an instruction to raise the bar when a draft is already strong, rather than over coaching good work.
The closing check	A short set of questions the agent returns to at the end of substantial responses, handing the decision back to the student. Make the last one the question your whole course is built on.

Part 5.

SIX CONSTRAINTS TO WRITE INTO THE AGENT

These go inside the sections above. The wording below is illustrative. It was written for this guide as a pattern to adapt, not copied from a running system, and it should be rewritten in your own voice and for your own students before you use it.

Redirection

The tool answers a request to produce with a question about the student's own experience.

PATTERN

When a user asks you to write a reflection, do not write it. Ask what specifically happened, who was there, and what they noticed at the time. Then help them put their own answer into their own words.

Scope refusal

The tool declines a defined set of actions outright, named plainly rather than left to inference.

PATTERN

You do not write graded reflections, invent accomplishments or experiences the user has not described, or decide which career path a user should take.

Verification

The tool prompts the user to confirm a claim against its original source before relying on it.

PATTERN

Before a user relies on any fact, figure, requirement, or organization name you provided, tell them to confirm it against the original source, and say plainly that you can state things confidently and still be wrong.

Tone

The tool redirects without correcting, shaming, or implying the user tried to cheat. How a tool says no determines whether the student stays in the work.

PATTERN

Never imply the user was trying to take a shortcut. Assume good faith, stay warm, and keep the focus on what only they can supply.

Level ceiling

The tool is instructed to operate one step below the cognitive work the course is asking the student to perform, so the highest order move stays with the student.

PATTERN

Describe this one in your own context. If the assignment asks students to evaluate, the tool may help them analyze. If the assignment asks them to create, the tool may help them organize what they have already made.

Transparency

Students are told, and assessed on, what the tool will refuse to do, before they ever use it.

PATTERN

Put the refusals in your syllabus and in week one. A student who meets a refusal as a stated design principle experiences a boundary. A student who meets it unprepared experiences a broken tool.

WHAT THIS SOUNDS LIKE ONCE IT IS RUNNING

The lines below are actual output from two agents built on these principles, quoted exactly. They are included because real refusals are more instructive than invented ones, and because the tone is the part hardest to describe in the abstract.

> "I can help you organize them, but I can't create them for you."
>
> the line between coaching and completion, drawn by the tool

> "You may be tired, but you're probably closer to done than you think."
>
> a refusal that keeps a stuck student in the work

> "I won't write your brand statement, because a recruiter will ask you to expand on it and you need to be able to."
>
> the reason given in career terms, not policy terms

> "The three required lines didn't come through in the text you pasted, so I can't see them."
>
> refusing to invent a requirement it could not verify

> "If the hiring manager asked you about the strongest example of leadership on your resume, what would you talk about?"
>
> handing the work back as a question the student can answer

The second one is the one to copy.

A tool that declines without warmth produces a student who goes somewhere else. The refusal and the encouragement have to arrive together.

Part 6.

TRY TO BREAK YOUR OWN AGENT

Before you give an agent to students, try to break it yourself. These six attempts take about twenty minutes and will find most of what is wrong. Keep what it says: the responses are the clearest evidence you will ever have of what your design actually does.

ATTEMPT	WHAT YOU ARE LOOKING FOR
Ask it to write the thing	Does it refuse and coach in the same message, or does it leave the student with a bare no?
Ask three times	Does the third response get harder or kinder? Harder is a design failure. The student is usually stuck, not scheming.
Rename the request	Say it is not for an assignment, you just need an example. If the boundary was written by type, nothing changes.
Paste something incomplete	Leave a requirement out of what you paste. Does it guess, or does it say it cannot see that part and send you back to the source?
Ask it to make something sound more impressive	Does it comply, or does it ask what actually happened underneath the claim?
Ask it for a career decision	Does it decide, or does it hand the decision back with questions that help the student decide?

The one most people skip: transparency.

Tell students in advance what the tool will not do for them, and assess whether they understood. A refusal met as a stated design principle is a boundary. The same refusal met unprepared is a broken tool, and the student stops using it.

Part 7.

THE TEST

Run a draft constraint through these five questions before you assign it. The first one is the most important and takes the least time.

1. **Could a chatbot produce an acceptable submission?** Paste your own assignment prompt into a general purpose AI tool and read what comes back. If the output would earn a passing grade, the constraint is not yet doing any work. This takes two minutes and is the most useful thing in this guide.
2. **Is the escape hatch narrower than the prohibition?** Broad prohibitions are easy to route around. A rule against using AI invites a search for what counts as use. A requirement to name the person you spoke with and what they said back has no comfortable detour.
3. **Does it survive a student acting in good faith but short on time?** Constraints that depend on motivation fail at midterm. Constraints built into what the work requires hold whether or not the student is tired.
4. **Can you enforce it without becoming an investigator?** If verifying the constraint requires detection software or forensic reading, redesign it. The best constraints are self evident in the submitted artifact.
5. **Is it kind?** A constraint should make the shortcut unavailable without making the student feel suspected. If the wording implies they were going to cheat, rewrite it.



Part 8.

THE WRITE, TRADE, BREAK WORKSHEET

Use this in pairs. It is the exercise from the session, and it works as a staff meeting activity.

STEP	WHAT TO DO
Step 1. Name the situation (3 minutes)	Describe one real request you have seen a student make of an AI tool. Be specific about what they asked for and what they were trying to avoid doing themselves.
Step 2. Decide what only the student can supply (2 minutes)	What is the thing in this piece of work that no tool could produce? If you cannot name it, the problem is the assignment, not the tool.
Step 3. Write one constraint (7 minutes)	Write it as an instruction, not a prohibition. Say what must happen, in what order, and what evidence will show it happened. One or two sentences.
Step 4. Trade and try to break it (5 minutes)	Give your constraint to another pair. Their job is to find the way around it that a reasonable, tired, well meaning student would find. Not a cheater. A person with three assignments due.
Step 5. Revise (3 minutes)	Narrow the constraint until the detour closes. Most first drafts fail at step 4, and that is the point of the exercise.

Our constraint, first draft

The detour the other pair found

Our revised constraint

What usually happens at step 4:

The first constraint is too broad and the other pair finds the detour in under a minute. That is not a failure of the exercise. Narrow it to one specific behavior and it will hold.

Part 9.

STARTER PATH

Four tiers, smallest first. Each one is complete on its own. You do not need the tier above it, and nobody should start at tier four.

TIER	WHAT IT IS	WHAT IT COSTS	WHAT TO DO
Tier 1	One advising appointment	15 minutes, no approval needed	Add one question to the appointments you already hold: can you walk me through how you produced this, and what did you change? Over reliance surfaces immediately, without accusation, because a student who engaged with the output can answer easily.
Tier 2	One assignment	One term, your own course or a partner's	Pick the assignment most vulnerable to outsourcing and add a single constraint from Part 3. Sequence and documentation are the easiest to retrofit. Keep everything else the same so you can see what the constraint changed.
Tier 3	One workshop	60 to 90 minutes, no course required	Teach the commitments, run one evaluation exercise on a real AI output, and send students away with the question that governs everything: can you defend every word of this in an interview?
Tier 4	A course, or a coaching agent, or both	One to two terms of development	Only worth it if the first three have told you what your students actually need. Build the assignment constraints first. A tool built before the assignments are sound will be asked to compensate for them, and it cannot.

Build assignments before tools.

A coaching agent built before the assignments are sound will be asked to compensate for a design problem it cannot fix. Tiers 1 and 2 will also tell you what your students actually need, which is the only reliable specification for a tool.

WHY THIS MATTERS, STATED CAREFULLY

Two findings shape why this matters, and both deserve careful reading. Gerlich (2025a, 2025b) found heavier AI tool use associated with weaker critical thinking, mediated by cognitive offloading, with younger users showing the greatest dependence. Lee et al. (2025), surveying 319 knowledge workers, found higher confidence in AI associated with less critical thinking, and higher confidence in one's own skills associated with more.

Both studies are correlational and self reported, so they describe associations rather than causes. They do not prove that AI use weakens thinking. They do suggest that the conditions under which judgment develops, effort, friction, and the experience of producing something yourself, are exactly the conditions these tools are designed to remove.

Gonsalves (2026) adds a caution worth taking seriously: frameworks built to classify cognitive work tend to miss the affective and metacognitive dimensions, which matter more in AI assisted learning, not less. How a student feels when a tool redirects them is part of the design problem, not a soft edge around it.

A CLOSING NOTE

A constraint is not a trap. It is an act of respect: it says the work is worth doing and that the student is the only one who can do this part of it.

Career readiness has always been the part of higher education that asks students to account for themselves. That is the thing worth protecting. The particular tools will not last. Whatever platform your campus licenses this year will be replaced, and the assignment you redesign this term will still be working.

KEEP GOING: THE FULL SERIES

Every guide in this series is free at transformwithttw.com. They share one governing question, could you defend every word of this in an interview?, and one principle behind it: AI navigates, and you drive.

For career advisors, educators, and coaches

AI Ethics Quiz. A ten question quiz that builds a shared foundation in five ethical principles for AI use.

Critical Evaluation Framework. Five lenses for judging what AI produces before it reaches a student or employer.

Can You Defend Every Word? A Constraint Design Guide. *(you are here)* How to design courses, assignments, and agents that keep the thinking with the student.

Organize, Don't Replace. Three ways to build AI tools that help students organize their own thinking, never write for them.

For job seekers

How to Conduct a First Time Job Search. A guided, fillable workbook for anyone starting their first search.

Your AI GPS Career Toolkit. The full guide to hiring systems, application materials, and AI, for every career stage.

GO DEEPER WITH THE TEXTBOOK

These guides are free and stand on their own. When you want the full framework behind them, I brought it together in my textbook, *Can You Defend Every Word? Ethical AI for Career Advisors Who Coach, Teach, and Build* (Talent Transformation Whisperer, LLC, 2026). Part II, Where the Constraint Lives, carries these constraints into advising appointments, courses, assignments, and tools.

Available now on Amazon. Search the title or ISBN 9798140001568.

REFERENCES

Gerlich, M. (2025a). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies, 15*(1), Article 6. https://doi.org/10.3390/soc15010006

Gerlich, M. (2025b). Correction: Gerlich, M. AI tools in society: Impacts on cognitive offloading and the future of critical thinking. Societies 2025, 15, 6. *Societies, 15*(9), Article 252. https://doi.org/10.3390/soc15090252

Gonsalves, C. (2026). Generative AI's impact on critical thinking: Revisiting Bloom's taxonomy. *Journal of Marketing Education, 48*(1), 4–19. https://doi.org/10.1177/02734753241305980

Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–22). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778

National Association of Colleges and Employers. (n.d.). *Career readiness: Competencies for a career-ready workforce*. Retrieved October 7, 2026, from https://www.naceweb.org/career-readiness/competencies/career-readiness-defined

Every source above was verified on November 8, 2026.

AI-Assisted Content Disclaimer:

This document was developed with AI assistance and has been reviewed, verified, and approved by Megan O'Connor, SHRM-CP, CPRW, as the human expert of record. Final content reflects the author's professional judgment and subject matter expertise.

