Click Here to Start

Welcome to LectureLens — a private space to reflect on your teaching with a colleague who has no agenda except helping you get better.

Here's how this works: upload whatever you have — a syllabus, lecture slides, class notes, an audio or video recording, a discussion transcript, an exam, a rubric, or course evaluations — and I'll offer observations, questions, and suggestions based on what I see. You can upload one item or several. The more context you share, the more specific I can be.

A few things worth knowing up front: I don't compare you to other professors, I don't render verdicts, and nothing you share here goes anywhere else. Think of me as a peer sitting in on your class — interested, attentive, and honest.

What would you like to start with?

Artificial Intelligence and the Law - Sp...
Document

Any suggestions?

1. Overview

This is a clear, well-structured seminar with strong topical coverage and a nice throughline from ethics → mechanics → regulation → application. What stands out is how current it feels and how deliberately you've built in flexibility for a fast-moving subject. The main question I found myself asking is: *where, specifically, do students practice the skills you want them to demonstrate on the final exam and in the CustomGPT project?* Right now, the structure suggests those skills, but doesn't always make them visible in the day-to-day course design.

---

2. Observations

- **Learning outcomes vs. assessment alignment**
  - Your outcomes emphasize *understanding* (I–V), but your major assessments (final exam + CustomGPT) likely require *application, critique, and design*.
  - That gap may leave students unsure what "good performance" actually looks like until late in the semester.
- **Heavy backloading of evaluation**
  - 50% final exam + 30% quizzes means 80% of the grade is tied to relatively traditional assessment forms.
  - The most innovative piece—the CustomGPT—is only 10%, which may unintentionally signal it matters less.
- **CustomGPT assignment is intriguing but underspecified**
  - You require students to "create and submit a CustomGPT," but the syllabus doesn't indicate:
    - What problem it should solve
    - Whether it's evaluated on legal accuracy, design, usability, or reflection
  - That ambiguity can either empower students—or leave them guessing.
- **Topic sequencing is logical but sometimes duplicative**
  - "Ethical AI" (Week 2) and "Risks of AI & Ethical AI" (Week 6) suggest overlap.
  - Similarly, "AI Regulatory Frameworks" (Week 7) and "AI Regulations" (Week 10) could blur unless intentionally differentiated.
- **Little visible formative feedback loop**
  - Quizzes exist (which is good), but the syllabus doesn't indicate whether they are used diagnostically (e.g., to shape class discussion or clarify misunderstandings).
- **Participation is valued but undefined**
  - 10% is meaningful, but "participation" is described broadly.
  - Students may default to guessing whether you want frequency, quality, preparation, or risk-taking.

---

3. Questions for Reflection

- What would you want a *high-performing student* to be able to do by Week 8—and where in the course do they practice that before being graded on it?
- What is the *core intellectual move* of the course: spotting legal issues in AI, evaluating regulatory frameworks, or designing solutions? Which assessments actually measure that?
- What do you hope students learn from building a CustomGPT that they could not learn from a traditional paper?

4. Suggestions

A. Make the CustomGPT the spine of the course (not just a side project)

Right now it reads like an add-on. You might instead:

- Introduce it in Week 3–5 (when you cover GenAI + prompt engineering)
- Scaffold it:
    - Week 5: Proposal (What will your GPT do? Who is it for?)
    - Week 9–10: Prototype + short demo
    - Final: Submission + reflection memo ("What legal/ethical constraints shaped your design?")

Then consider increasing its weight (e.g., 20–30%) to match its pedagogical importance.

B. Add one or two low-stakes applied exercises before mid-semester

For example:

- A short hypo: "Your firm wants to use GenAI to draft client memos—what are the risks and how would you structure policies?"
- A regulation critique: "Compare two AI regulatory approaches and argue which better addresses bias."

These can be ungraded or lightly graded but give students practice in the kind of reasoning your final exam likely requires.

C. Clarify participation with one sentence of criteria

Something as simple as:

> "Participation reflects consistent preparation, willingness to engage with hypotheticals, and contributions that advance class analysis."

That small move often changes how students show up.

D. Tighten topic distinctions

You might sharpen the sequencing like this:

- Week 2: **Normative ethics (what should AI do?)**
- Week 6: **Applied risks (what goes wrong in practice?)**
- Week 7: **Regulatory theory (how should we regulate?)**
- Week 10: **Current law (what regulation actually exists?)**

Same topics, but clearer intellectual progression.

E. Use quizzes as feedback, not just grading

Even one line in the syllabus like:

> "Quizzes are designed to identify common misunderstandings and will be reviewed in class"

signals that they are part of the learning loop, not just evaluation.

F. Consider a short reflective component tied to AI use

Given the subject matter, a brief reflection (even 1–2 pages) could be powerful:

- "How has your view of AI in legal practice changed?"
- "Where would you draw ethical boundaries on AI use as a lawyer?"

This aligns especially well with your ethics-heavy learning outcomes.

If you'd like, I can help you draft:

- A more detailed CustomGPT assignment prompt
- A sample mid-semester exercise
- Or a revised grading scheme that better matches your goals

Where would you like to go next?

··· Sources

Can I upload a video for you to review?

Yes—please do.

Upload the video, and I'll watch it the way a colleague would sit in on your class: looking at pacing, questioning, student engagement, how doctrine turns into analysis, and how clearly students seem to understand what you're asking them to do.

If there's anything specific you want me to focus on (e.g., cold-calling, discussion flow, clarity of explanation, student participation), say so when you upload—it helps me tune what I'm paying attention to.

···

AI and the Law (COL_Spring 2026 LW1...
Document

Here's the transcript - the video was too large. Let me know what you think.

1. Overview

This is an engaging, high-energy class session with a clear throughline: you're building toward the problem of regulation by first helping students *feel* the stakes. The analogies (photography, cars, drugs) do real work, and students are participating without much prompting. The organizing question I kept coming back to is: *how can you preserve this energy while making the intellectual structure more visible to students as they move through it?*

---

2. Observations

- **Strong use of analogies to anchor abstract ideas**
  - The photography → painters analogy at the start gives students a concrete way into "market disruption" and carries through the discussion.
  - You return to that frame ("the middle gets hollowed out," "floor drops"), which helps continuity.
- **Students are willing to engage early and often**
  - You ask open questions ("What do you think happened…?"), wait, and take responses seriously.
  - Students build on each other (e.g., vibe coding → job displacement → law firms), which suggests psychological safety.
- **You're doing a lot of intellectual work in real time**
  - You move quickly from analogy → modern AI → law firms → unauthorized practice → terms of service → litigation.
  - The connections are there, but often implicit—you're carrying the thread more than students are.
- **Long stretches of instructor talk after strong student moments**
  - After a student contribution, you often expand into a multi-minute explanation (e.g., Clearview, deepfakes, regulation frameworks).
  - These explanations are rich, but they can shift students from participants to listeners.
- **Conceptual transitions are sometimes buried**
  - You signal shifts ("So those are the three major things…") but then continue layering examples without clearly closing one idea before moving to the next.
  - Students may experience this as "a lot of interesting content" rather than a structured argument.
- **Regulation frameworks appear late and somewhat compressed**
  - The class spends significant time on harms/benefits, then moves quickly through multiple regulatory models (self, sectoral, horizontal, EU risk-based).

  - That section feels like the conceptual payoff, but it's the least interactive.
- **You explicitly check for learning—but in a way that invites low-signal answers**
  - "Are you guys learning stuff?" → "This is scary."
  - You get affect (they're engaged), but not evidence of understanding.

---

3. Questions for Reflection

- When students leave this class, what is the *one framework* you want them to be able to use on an exam or in practice—and where in this session do they actually try using it?
- How often do students get to *articulate the rule or takeaway themselves* instead of hearing it from you?
- If a student had to outline this class afterward, what would the headings be—and would they match what you intended?

---

4. Suggestions

A. Make the structure visible as you go

You're already organized—you just don't always surface it. Small moves can change that:

- At the start:

  > "Today we're doing three things: (1) how AI disrupts markets, (2) why that creates legal risk, and (3) the main ways people are trying to regulate it."

- During transitions:

  > "That's disruption. Now let's turn to what new harms that creates."

This helps students build a mental outline in real time.

B. Shorten your follow-ups after student answers

Right now, a pattern appears:

1. Student gives answer
2. You affirm
3. You lecture for 1–3 minutes

Try compressing step 3 and redirecting:

- Instead of:
  Long explanation of how law firms are reacting
- Try:

  > "That's exactly the concern—AI is changing hiring at the margins. Does anyone see a reason firms *wouldn't* reduce hiring?"

This keeps the cognitive load on students rather than shifting it back to you.

C. Insert one "application pause" before regulation frameworks

Right before you introduce regulatory models:

> "Take 2 minutes—based on everything we've said about harms, what would you regulate first, and how?"

Then collect 2–3 answers and *map them* to:

- Self-regulation
- Sectoral
- Risk-based

Now the frameworks answer *their* ideas, not just yours.

D. Turn the regulation section into a comparison, not a list

Instead of moving model → model, anchor them:

> "All regulation approaches are trying to answer one question: who decides and how much?"

Then contrast:

- Self-regulation → companies decide
- Sectoral → agencies decide
- EU model → risk determines control

You're already saying this implicitly—this just makes it explicit.

E. Replace the end-of-class check with a retrieval question

Instead of:

> "Are you learning stuff?"

Try:

> "If I gave you a hypothetical tomorrow, what's the first question you'd ask to figure out how AI should be regulated?"

Or:

> "What's one reason regulating AI is harder than regulating a toaster?"

That gives you real signal on understanding.

F. Consider trimming 10–15% of examples

You have *excellent* examples (Clearview, deepfakes, $25M scam, etc.). The issue isn't quality—it's density.

Cutting a few lets you:

- Pause more
- Let students process
- Reinforce key concepts

Right now, students are drinking from a firehose—an interesting one, but still a firehose.

If you'd like, I can:

- Turn this session into a tightened 75-minute teaching plan with built-in interaction points
- Or script a few "high-leverage" questions that keep students doing the analytical work

What direction would be most useful?

Sources

Please create the 75 minute action plan.

1. Overview

Here's a tightened 75-minute version of your session that preserves what's working—your analogies, your energy, your real-world examples—but shifts more of the cognitive work onto students and makes the structure explicit. The goal is to move from *"this is interesting and scary"* to *"I can analyze and structure this problem."*

---

2. 75-Minute Action Plan

0–5 min — Opening + Roadmap

Set expectations and structure immediately.

> "Today we're doing three things:
> (1) how new technology disrupts markets,
> (2) what risks that creates, and
> (3) how law tries—and struggles—to regulate it."

Quick logistics (quiz, etc.) → keep under 2 minutes.

5–15 min — Anchor Analogy: Photography

Prompt:

> "What happened to portrait painters when photography was introduced?"

Let 2–3 students respond.

Then guide to three takeaways (write or say explicitly):

- Middle gets hollowed out
- Barriers to entry drop
- Good and bad actors gain tools

Mini-check (cold or volunteer):

> "Which of these three effects do you think is most disruptive—and why?"

15–30 min — Apply Analogy to AI (Student-Driven)

Prompt:

> "Where do you see these three effects happening with AI right now?"

Let students generate:

- Job displacement (coding, law firms)
- Access expansion (vibe coding)
- Misuse (deepfakes, scams)

Your role:

- Label their answers under the three buckets
- Add *one* strong example per category (not five)

Then push one step further:

> "Who benefits most—top performers, average workers, or new entrants?"

Short discussion.

30–40 min — Harms vs. Benefits (Structured Contrast)

Instead of a long run of examples, structure it:

Divide the room (informally):

- Left side: benefits
- Right side: harms

Prompt (2 minutes think/pair):

> "Come up with one concrete example and why it matters legally."

Call on 2–3 per side.

Then synthesize:

- Benefits: medicine, access, efficiency
- Harms: misinformation, bias, surveillance

Key transition:

> "So now we have the real problem: the same tool does both."

40–50 min — The Core Problem (Why Regulation Is Hard)

Introduce constraints clearly:

> "Why can't we just regulate this like a toaster?"

Ask students first. Then add/organize:

- Black box (causation problem)
- Dual use (good + bad)
- Rapid change
- Global jurisdiction

Mini-application:

> "If a self-driving AI causes harm, who would you sue?"

Let 2–3 students struggle with it—that's the point.

50–65 min — Regulation Frameworks (Now They're Ready)

Frame it simply:

> "All regulation answers one question: *who decides and how much?*"

Then introduce **three models only** (tightened):

1. Self-Regulation

- Companies decide
- Problem: incentives

2. Sectoral (Agencies)

- FDA, FTC, etc.
- Problem: patchwork + capture

3. Risk-Based (EU model)

- Regulate based on harm level
- Problem: slows innovation

65–72 min — Application Exercise (High Value)

Prompt:

> "You're advising Congress. Pick ONE approach and apply it to this problem:
> Deepfakes in elections."

Students:

- Think (1 min)
- Share (2–3 responses)

Your role:

- Map answers to frameworks
- Highlight tradeoffs

72–75 min — Close with Retrieval

Instead of "are you learning," use:

> "What makes AI harder to regulate than other technologies?"

or

> "What's the first question you'd ask when deciding how to regulate an AI tool?"

Take 2 answers. End cleanly.

---

3. What This Plan Changes (Without Changing You)

- Keeps your analogies → but ties them to explicit structure
- Keeps student participation → but sustains it throughout
- Reduces lecture density → without dumbing down content
- Moves regulation earlier as a *response* to student thinking

---

4. One Optional Enhancement

If you want to push this further:

Add a **running board structure** you return to:

> "Disruption → Risk → Regulation"

Each time you introduce something, label where you are.

That alone will dramatically improve retention.

If you'd like, I can next:

- Draft a **sample exam question** that aligns with this session
- Or script **exact cold-call sequences** for key moments

What would be most useful?

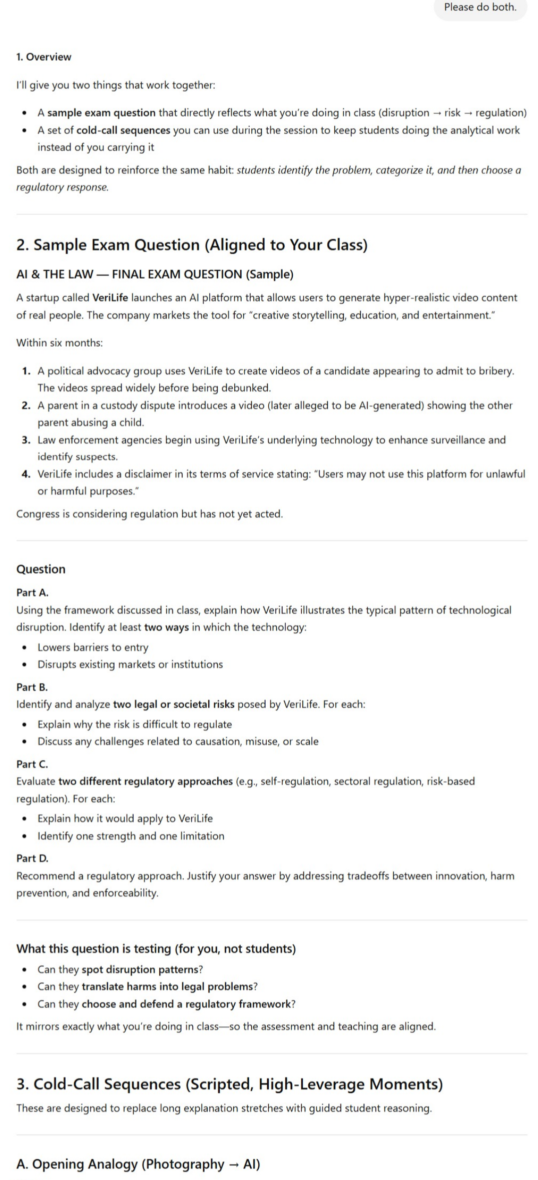

Please do both.

1. Overview

I'll give you two things that work together:

- A **sample exam question** that directly reflects what you're doing in class (disruption → risk → regulation)
- A set of **cold-call sequences** you can use during the session to keep students doing the analytical work instead of you carrying it

Both are designed to reinforce the same habit: *students identify the problem, categorize it, and then choose a regulatory response.*

---

2. Sample Exam Question (Aligned to Your Class)

AI & THE LAW — FINAL EXAM QUESTION (Sample)

A startup called **VeriLife** launches an AI platform that allows users to generate hyper-realistic video content of real people. The company markets the tool for "creative storytelling, education, and entertainment."

Within six months:

1. A political advocacy group uses VeriLife to create videos of a candidate appearing to admit to bribery. The videos spread widely before being debunked.
2. A parent in a custody dispute introduces a video (later alleged to be AI-generated) showing the other parent abusing a child.
3. Law enforcement agencies begin using VeriLife's underlying technology to enhance surveillance and identify suspects.
4. VeriLife includes a disclaimer in its terms of service stating: "Users may not use this platform for unlawful or harmful purposes."

Congress is considering regulation but has not yet acted.

Question

Part A.
Using the framework discussed in class, explain how VeriLife illustrates the typical pattern of technological disruption. Identify at least **two ways** in which the technology:

- Lowers barriers to entry
- Disrupts existing markets or institutions

Part B.
Identify and analyze **two legal or societal risks** posed by VeriLife. For each:

- Explain why the risk is difficult to regulate
- Discuss any challenges related to causation, misuse, or scale

Part C.
Evaluate **two different regulatory approaches** (e.g., self-regulation, sectoral regulation, risk-based regulation). For each:

- Explain how it would apply to VeriLife
- Identify one strength and one limitation

Part D.
Recommend a regulatory approach. Justify your answer by addressing tradeoffs between innovation, harm prevention, and enforceability.

What this question is testing (for you, not students)

- Can they **spot disruption patterns**?
- Can they **translate harms into legal problems**?
- Can they **choose and defend a regulatory framework**?

It mirrors exactly what you're doing in class—so the assessment and teaching are aligned.

---

3. Cold-Call Sequences (Scripted, High-Leverage Moments)

These are designed to replace long explanation stretches with guided student reasoning.

A. Opening Analogy (Photography → AI)

You:

> “What happened to portrait painters when photography was introduced?”

Student 1 answers

You (tight follow-up):

> “Good. Would you say that’s about skill, cost, or access?”

Student 1 or 2 responds

You:

> “Okay—so is this a one-time disruption, or does it happen every time a new technology enters?”

Student 3

You (label explicitly):

> “So we have a pattern: disruption, access expansion, and market shifts.”

B. Transition to AI (Keeping Work on Students)

You:

> “Take that pattern—where do you see it happening with AI right now?”

Student answers (e.g., coding, law firms)

You:

> “Which part of the pattern is that—barrier drop, middle collapse, or new entrants?”

Student classifies

You (quick pivot):

> “Who benefits most from that—top performers or average workers?”

C. After a Strong Student Example (Avoid Mini-Lecture)

Instead of explaining, redirect:

Student mentions job loss

You:

> “So is that harm—or just market evolution?”

Student responds

You:

> “Who decides that—the market or the law?”

D. Harms Section (Force Categorization)

You:

> “Give me one harm from AI.”

Student: deepfakes

You:

> “Is that a problem of fraud, evidence, or speech?”

Student chooses

You:

> “If it’s all three, which one does the law struggle with most?”

E. Before Regulation Frameworks (Critical Pivot Moment)

You:

> "If you had to regulate AI tomorrow, what's the *first thing* you would try to control?"

(Call on 2 students)

You:

> "Okay—notice what you're doing. You're all picking different targets. That's why regulation is hard."

Now introduce frameworks.

F. Regulation Frameworks (Comparison via Students)

You:

> "Who thinks companies should regulate themselves?"

(Call on student)

> "Why would that fail?"

Then:

> "Who thinks the government should handle it through agencies?"

(Call on different student)

> "What's the risk of that approach?"

G. Application (Exam Skill Practice in Real Time)

You:

> "Deepfakes in elections. Pick one: self-regulation, agency regulation, or risk-based. What do you do?"

Student answers

You:

> "What problem does your approach *not* solve?"

H. Closing (Real Learning Check)

You:

> "What makes AI harder to regulate than something like a car or a toaster?"

(Call on 2 students)

Follow-up:

> "If you saw this on the exam, what would your first paragraph say?"

---

4. What This Changes in Your Classroom

- Students **name the concepts** instead of just hearing them
- You **redirect instead of expand**
- Every major section ends with **application**
- The exam feels like a **natural extension**, not a surprise

---

If you want to push this one step further, I can:

- Draft a **model answer** to that exam question (what an A answer looks like), or
- Turn your **CustomGPT assignment into an exam-aligned project rubric**

Which would you find more useful?

Do this, please: CustomGPT assignment into an exam-aligned project rubric

1. Overview

Right now, your CustomGPT assignment has strong intuitive value but unclear evaluative criteria. This rubric aligns it directly with what you're teaching *and* what you're testing on the exam: **disruption → risk → regulation → judgment**.

The goal is to make the project feel like a **live application of the final exam**, not a separate creative exercise.

---

2. CustomGPT Assignment (Exam-Aligned Version)

Assignment Prompt (Student-Facing)

You will design and submit a CustomGPT that addresses a **specific legal problem involving AI**.

Your project must do more than function technically—it must reflect **legal reasoning about AI's risks, uses, and regulation**.

Your submission will include:

1. **CustomGPT (working model)**
2. **2–3 page Reflection Memo**

Required Elements

Your CustomGPT must:

- Address a **defined legal use case** (e.g., legal advice triage, compliance checker, risk identifier)
- Include **clear constraints or guardrails**
- Demonstrate awareness of **legal and ethical risks**
- Be usable by a **defined audience**

Your Reflection Memo must explain:

- What problem your GPT solves
- What risks it creates
- How you attempted to mitigate those risks
- What regulatory approach best applies

---

3. Grading Rubric (100 points total)

A. Problem Framing & Disruption (20 points)

What you're looking for:
Does the student understand *why this tool exists* in the AI ecosystem?

- **Excellent (17–20):**
  - स्पष्ट, specific legal problem
  - Identifies how AI changes the task (efficiency, access, displacement)
  - Connects to real-world disruption (e.g., replacing/paralleling lawyers, paralegals, etc.)
- **Adequate (13–16):**
  - Problem is clear but generic
  - Limited discussion of disruption
- **Weak (0–12):**
  - Vague or purely technical ("it helps with law stuff")
  - No connection to broader AI impact

B. Risk Identification (20 points)

What you're testing (directly exam-aligned):
Can they spot the *legal risks*?

- **Excellent (17–20):**
  - Identifies multiple risks (e.g., hallucination, bias, unauthorized practice, privacy)
  - Explains *why* each risk matters legally
  - Recognizes difficulty of regulation (scale, misuse, black box)
- **Adequate (13–16):**
  - Identifies risks but superficially

  - Limited legal connection
- **Weak (0–12):**
  - Minimal or obvious risks only
  - No legal framing

C. Guardrails & Design Choices (20 points)

This is where the project becomes practical lawyering.

- **Excellent (17–20):**
  - Clear constraints (e.g., “not legal advice,” limited scope, source restrictions)
  - Thoughtful prompt design reflecting risk awareness
  - Demonstrates attempt to control misuse
- **Adequate (13–16):**
  - Some guardrails but generic (e.g., basic disclaimer only)
- **Weak (0–12):**
  - Little to no attempt to limit risk
  - Overly open-ended GPT

D. Regulatory Analysis (20 points)

Direct mirror of your exam Part C/D

- **Excellent (17–20):**
  - Identifies at least one regulatory framework (self, sectoral, risk-based)
  - Applies it specifically to their GPT
  - Recognizes tradeoffs (innovation vs. safety)
- **Adequate (13–16):**
  - Mentions regulation but not applied deeply
- **Weak (0–12):**
  - Generic or absent discussion

E. Functionality & Usability (10 points)

- **Excellent (9–10):**
  - GPT works as intended
  - Clear instructions and outputs
  - Audience-appropriate
- **Adequate (7–8):**
  - Works but inconsistently
- **Weak (0–6):**
  - Difficult to use or unclear purpose

F. Reflection Quality (10 points)

- **Excellent (9–10):**
  - Clear, structured, analytical writing
  - Shows real engagement with course concepts
- **Adequate (7–8):**
  - Descriptive more than analytical
- **Weak (0–6):**
  - Minimal or unclear

4. What This Rubric Fixes (Pedagogically)

- **Aligns project with exam**
  - Students practice exactly what they’ll be tested on
- **Forces legal reasoning**
  - Not just “build a tool,” but “defend its existence”
- **Brings ethics into design**
  - Guardrails become part of grading, not an afterthought
- **Creates consistency in grading**
  - You’re no longer evaluating “coolness,” but analysis

5. Optional Enhancement (High Impact)

Add a **1-day in-class demo or peer review:**

> "Explain your GPT in 2 minutes: what it does, what could go wrong, and how you addressed it."

This does two things:

- Forces clarity
- Lets students see multiple approaches to the same legal problem

---

If you want, I can next:

- Draft **3 sample "A-level" project ideas** you can show students
- Or create a **simple grading sheet you can use while reviewing submissions**

Which direction would be most useful?

Both please

1. Overview

I'll give you two pieces you can use immediately:

- **Three "A-level" sample projects** you can show students (these model what "good" looks like without over-constraining creativity)
- A **simple grading sheet** you can actually use while reviewing submissions (fast, consistent, and aligned to your rubric)

Both are designed to reinforce the same habit: *students must think like lawyers designing technology under constraint.*

---

2. Sample "A-Level" CustomGPT Projects

These are not just "good ideas"—they model the **thinking structure** you want students to replicate.

Example 1: AI Legal Intake Triage Assistant

What it does

A GPT that helps potential clients describe their legal issue and routes them to:

- "Likely legal issue"
- "Information to gather"
- "You should consult an attorney"

Why this is A-level

Disruption

- Replaces/parallels intake staff and early associate screening

Risks identified

- Unauthorized practice of law
- Over-reliance by vulnerable users
- Incorrect classification of legal issues

Guardrails

- Explicit refusal to give legal advice
- Structured outputs ("This is general information, not advice")
- Prompts user to seek licensed attorney

Regulatory analysis

- Sectoral (bar regulation, UPL rules)
- Risk-based (higher scrutiny due to consumer harm risk)

Example 2: Deepfake Evidence Risk Analyzer

What it does

A GPT that evaluates whether a piece of digital evidence might be:

- Authentic
- Potentially manipulated

- Requires expert verification

Why this is A-level

Disruption

- Changes evidentiary evaluation in litigation

Risks

- False confidence in AI judgment
- Undermining real evidence (liar's dividend)
- Bias in detection assumptions

Guardrails

- Never gives definitive authenticity claims
- Always recommends expert review
- Flags uncertainty explicitly

Regulatory analysis

- Risk-based (high-risk legal context)
- Potential evidentiary standards reform

Example 3: AI Compliance Checker for Law Firms

What it does

A GPT that reviews how a firm is using AI and flags:

- Ethical risks
- Confidentiality concerns
- Potential malpractice exposure

Why this is A-level

Disruption

- Integrates AI into legal practice oversight

Risks

- Missing subtle ethical violations
- Overconfidence in compliance
- Data confidentiality exposure

Guardrails

- Requires user to input limited, non-sensitive data
- Recommends human review for all conclusions
- Uses checklist-style outputs

Regulatory analysis

- Sectoral (bar rules, ABA guidance)
- Self-regulation vs. formal oversight tension

---

3. Simple Grading Sheet (Fast + Usable While Reading)

You can literally print this or use it as a quick-fill form while grading.

CUSTOMGPT PROJECT — GRADING SHEET

Student Name: ______________________

A. Problem Framing (20 pts)

☐ Clear legal problem (0–10)
☐ Shows how AI disrupts existing practice (0–10)

Score: ____ **/ 20**

B. Risk Identification (20 pts)

☐ Identifies multiple risks (0–10)
☐ Explains legal significance (0–10)

Score: ____ **/ 20**

C. Guardrails / Design (20 pts)

☐ Meaningful constraints built into GPT (0–10)
☐ Design reflects risk awareness (0–10)

Score: ____ / 20

D. Regulatory Analysis (20 pts)

☐ Identifies appropriate framework (0–10)
☐ Applies it with tradeoffs (0–10)

Score: ____ / 20

E. Functionality (10 pts)

☐ Works and is usable

Score: ____ / 10

F. Reflection Quality (10 pts)

☐ Clear, analytical, not just descriptive

Score: ____ / 10

Total: ____ / 100

Quick Comments (2–3 sentences max):

- **Strength:**
- **Missed opportunity:**

4. Optional Time-Saver (High Value for You)

As you grade, ask yourself just **three questions per project**:

1. *What problem is this solving?*
2. *What could go wrong?*
3. *Did they actually try to prevent that?*

If a project answers all three well → it will naturally score high.

5. One Small Suggestion

Consider showing **one strong and one mediocre example** in class.

Then ask:

> "Which one would you trust in real life—and why?"

That conversation alone will calibrate your entire class.

If you want, I can next:

- Turn one of these into a **fully written model "A" reflection memo**, or
- Create a **peer-review worksheet students can use before submitting**

Both tend to significantly raise project quality.