

## Teacher Guide: Create: A Critical Policy Analysis of an AI Position Paper

This guide provides expected responses for each step, common student pitfalls, feedback strategies, and a full annotated example of a completed artifact (Step 7).

---

### Step 1 — Summarize the Core Claim

#### What a strong response looks like:

Students should capture three elements: (1) the prediction that AI will reach superintelligence within a few years, outperforming humans even when humans are AI-assisted; (2) the comparison to electricity, mass production, or something even larger; and (3) the call on governments, companies, and workers to enact policies now for a smoother transition.

#### Example strong response:

*“OpenAI predicts that artificial intelligence will reach 'superintelligence' within a few years — a point at which AI outperforms humans at virtually every task, even with human-AI collaboration. They compare this shift to transformative historical technologies like electricity and mass production, suggesting its impact could be even greater. To avoid severe societal disruption, they are urging governments, companies, and workers to begin implementing policies immediately.”*

#### Common pitfalls:

- Students may omit the call to action and only summarize the prediction.
- Students may confuse "superintelligence" with general AI improvement — emphasize the specific definition: outperforming humans *even when humans use AI tools themselves*.

**Feedback tip:** Ask students: "Does your summary tell me *what* OpenAI predicts, *how big* they think it is, and *what they want done about it*?"

---

## Step 2 — Identify the Historical Analogies

### What a strong response looks like:

- **The Industrial Revolution:** Used to illustrate that major technological shifts can cause widespread societal disruption and require deliberate political responses to manage them.
- **The New Deal:** Used as an example of a successful, proactive policy response that modernized the social contract during a period of economic upheaval — offered as a model for how government can respond to AI disruption.

### Common pitfalls:

- Students may list the events without explaining *why* they are being used as analogies (i.e., what argumentative work they do in the paper).
- Students may conflate the two — the Industrial Revolution is the *problem* analogy; the New Deal is the *solution* analogy.

**Feedback tip:** Ask: "Is this example being used to show that disruption happens, or to show that policy can help?"

---

## Step 3 — List and Explain the Five Negative Externalities

### Expected list (all five must be present):

1. **Job losses** — Superintelligent AI could automate vast categories of human work, leaving many workers unemployed with few alternatives.
2. **Further concentration of wealth** — The productivity and profit gains from AI are likely to flow primarily to AI companies and wealthy investors, widening inequality.
3. **Misuse by bad actors** — Powerful AI tools could be exploited by criminals, hostile states, or other bad actors to cause harm at unprecedented scale.
4. **Erosion of the personal income tax base** — As AI replaces human workers, the government collects less income tax revenue, threatening the funding of public services.
5. **Data centers competing with communities for resources** — Large AI data centers consume enormous amounts of electricity and water, potentially straining local infrastructure and resources.

### Common pitfalls:

- Students frequently omit the tax base erosion or the data center resource competition — these are less intuitive and should be flagged.

- Explanations may be too vague (e.g., "AI could be bad for jobs" without explaining *why* that is a problem).

**Feedback tip:** Prompt students to think about *who specifically is harmed* by each externality.

---

#### Step 4 — Evaluate the Proposed Solutions

##### What a strong response looks like:

Students should accurately identify the funding/solution mechanism and provide a reasoned personal assessment. Mechanisms mentioned in the video include:

- Tax increases on corporations and high earners
- Government funding of new programs
- AI companies contributing to a public wealth fund (described as a one-time seed contribution of an undetermined amount)

##### Example strong response (for "job losses"):

*“The paper proposes addressing job losses through government-funded retraining programs and a public wealth fund seeded by AI companies. While retraining programs have historical precedent, their effectiveness depends heavily on scale and political will — both of which are uncertain in the current U.S. political climate. The public wealth fund is particularly speculative since the contribution amount is undetermined, making it difficult to evaluate as a serious solution.”*

##### Common pitfalls:

- Students may accept the solutions uncritically without assessing feasibility.
- Students may incorrectly assign the wrong solution to the wrong externality.

**Feedback tip:** Ask students: "What would need to be true for this solution to actually work?"

---

#### Step 5 — Analyze the Source: Identify Potential Bias

**What a strong response looks like:**

Students should identify that as a for-profit company, OpenAI has a financial interest in avoiding product-specific regulations (taxes on AI products, usage restrictions) that would raise prices, lower demand, and cut revenue. The video compares this to oil and tobacco companies, which faced exactly this kind of product-specific regulation. Students should take a clear, reasoned position on whether this bias invalidates the argument.

**Example strong response:**

*“As a for-profit company, OpenAI benefits financially from avoiding product-specific regulations — the kind of taxes and restrictions that governments have imposed on harmful industries like oil and tobacco. Those regulations raise prices and reduce demand, directly cutting into company revenue. OpenAI’s advocacy for social policies rather than product regulations may therefore reflect its business interests as much as its stated concern for society. That said, this does not necessarily mean the paper’s analysis of risks is wrong — a company can identify real problems while simultaneously steering toward self-serving solutions.”*

**Common pitfalls:**

- Students may conclude that bias completely invalidates the paper (overcorrection) or dismiss the bias concern entirely (undercorrection).
- Encourage nuanced positions: bias shapes *which solutions are proposed*, not necessarily *whether the problems are real*.

---

**Step 6 — Distinguish Between Social Policy and Product Regulation****What a strong response looks like:**

Students should clearly explain that product-specific regulation (like a tax on AI services or a cap on AI deployment) would directly reduce OpenAI’s revenue and market share — similar to how tobacco taxes reduced cigarette sales and oil regulations raised costs for energy companies. Social policy (worker retraining, wealth funds, income support) addresses the *consequences* of AI without restricting the *product itself*. Therefore, it is in OpenAI’s interest to frame the conversation around social policy.

**Common pitfalls:**

- Students may understand the distinction conceptually but fail to connect it to OpenAI's incentives.
- Students may not use the industry comparison as instructed.

**Feedback tip:** Ask: "If the government taxed each AI query at \$0.01, how would that affect OpenAI differently than a worker retraining fund?"

---

## Step 7 — Full Critical Analysis

### What a strong response looks like:

A polished, 3–4 paragraph essay that synthesizes all prior steps in flowing prose. Paragraphs should move logically: summary → externalities and solutions → critical evaluation of motives → verdict.

### Common pitfalls:

- Reverting to bullet points or lists instead of prose.
- Writing a summary with no critical layer — students must include the bias/motive analysis.
- Vague concluding verdicts ("It's a mix of both") without supporting reasoning.

**Feedback tip:** Ask: "Could someone who didn't watch the video read your analysis and understand both what OpenAI argues *and* why a critical reader might be skeptical?"

---

## Full Annotated Example — Completed Artifact (Step 7)

---

### *"Critical Analysis of OpenAI's April 2026 Policy Paper*

*On April 6, 2026, OpenAI released a paper predicting that artificial intelligence will reach "superintelligence" within a few years — a threshold at which AI outperforms humans at virtually every task, even when humans themselves are using AI. The company compares the potential impact of this technology to electricity or mass production, arguing it could be transformative on a civilizational scale. Drawing on historical precedents like the Industrial Revolution and the New Deal, the paper argues that market forces alone will be insufficient to manage the disruption ahead,*

*and it calls on governments, companies, and workers to begin enacting policies now.*

*The paper identifies five specific negative externalities of superintelligence: widespread job losses, further concentration of wealth among the already-rich, misuse of AI by bad actors, erosion of the personal income tax base as human workers are replaced, and data centers competing with local communities for energy and water. For each problem, the paper proposes solutions funded through a combination of corporate and high-income tax increases, new government programs, and a one-time seed contribution by AI companies to a public wealth fund. While these mechanisms are reasonable in theory, their feasibility varies significantly — particularly in a political environment where major tax increases and new public funds face steep opposition.*

*However, it is worth considering who authored this paper and why. OpenAI is a for-profit company with a direct financial stake in how AI is regulated. In industries like oil and tobacco, governments have responded to major negative externalities with product-specific taxes and usage regulations that raise prices, reduce demand, and cut directly into company revenue. A paper that steers the policy conversation toward broad social programs — rather than inviting direct regulation of AI products — conveniently sidesteps that risk entirely. This does not mean the five externalities the paper identifies are fabricated; the risks are real and well-documented. But the selection of which solutions to propose appears shaped, at least in part, by OpenAI's business interests.*

*Ultimately, OpenAI's paper is best understood as something in between a genuine public service and a self-serving document. The company deserves credit for publicly acknowledging the downsides of the technology it is building and for calling for coordinated action. But readers should remain aware that the solutions it recommends are notably favorable to its own continued growth and freedom from product-specific constraints. Critical engagement with the paper — neither dismissing it outright nor accepting it uncritically — is the most productive approach.”*

---

**Annotation notes for teachers:**

- *Paragraph 1* accurately summarizes the claim, the scale comparison, and the historical framing — see Steps 1–2.

- *Paragraph 2* lists all five externalities correctly and notes the funding mechanisms while offering a feasibility caveat — see Steps 3–4.
- *Paragraph 3* explicitly uses the oil/tobacco comparison to analyze OpenAI's motives and distinguishes social policy from product regulation — see Steps 5–6.
- *Paragraph 4* offers a nuanced verdict that avoids overcorrection in either direction.

**Grade-level note:** This example represents a high-quality response. An acceptable response may be less sophisticated in its verdict paragraph but should demonstrate accurate recall of all five externalities and a clear articulation of the bias concern.