

EVERYDAY DIGITAL ETHICS · SEMINAR 02

CASE PHILOSOPHY SEMINAR

When Does AI-Assisted Work Stop Being Yours?

Achievement, authorship, learning, and human agency

PUBLIC-FACING QUESTION

If you can explain and defend an AI-written essay, is it still your work?

AI can now perform parts of thinking and writing that used to reveal a student's skill. The difficult question is not simply whether AI was used, but what the human must still contribute for the result to count as their achievement.

In one sentence	A seminar on whether authorship depends on labor, intellectual control, understanding, or responsibility - and whether an AI-supported ability can genuinely belong to a person.		
Preparation	Live session	Entry level	
35-45 minutes	About 55 minutes	No prior philosophy	

Format

Read the case, watch Material 0, then work through four guided texts. Selection priority: argument-bearing passages and objections, not abstracts or roadmap prose.

- Student-relevant
- Classic + current scholarship
- 4 texts + 1 video
- No policy debate required

01 Core case and three questions

Case: Sophie's excellent essay

Sophie is 17 and has to write a philosophy essay. She chooses the thesis herself and spends two hours arguing with an AI about objections. She rejects several suggestions, asks the AI to produce a structure, then asks it to draft large sections. She rewrites some paragraphs, checks every citation, and removes claims she cannot defend.

The final essay is far better organized than anything she could produce alone. In an oral conversation, Sophie can explain the thesis, answer objections, and say why she rejected alternatives. But if the same task were given tomorrow without AI, she could not produce an essay of comparable quality from a blank page.

Her teacher accepts that she understands the final argument but asks a different question: "Is this your achievement?" Sophie replies that architects do not make every brick and scientists use instruments; perhaps intellectual work can also be genuinely hers even when the tool does much of the execution.

Questions everyone answers

01	What must Sophie contribute for the essay to count as her work: effort, ideas, control over the process, understanding, or responsibility for the result?
02	If Sophie can perform the task reliably with AI but not without it, does she possess the relevant ability - or only access to an external ability?
03	If using AI produces a better essay but removes practice that would develop her judgment, when is efficiency a loss rather than a gain?

formation through practice -> authorship and epistemic control -> extended or curated ability

Opening poll: Clearly Sophie's work · Partly Sophie's work · Not Sophie's achievement · Depends on what she can explain and control

MATERIAL

0

Video Primer: The Extended Mind

User-selected YouTube primer · watch before the readings · directly relevant to Question 2

Question 2 asks whether Sophie can genuinely possess an ability that she cannot exercise without an external system. The extended-mind debate makes one assumption immediately contestable: a capacity does not have to be located entirely inside the individual in order to count as part of her cognition.

Why this belongs before the texts

Clark and Chalmers' challenge is functional rather than anatomical. If an external resource is tightly integrated into a person's ordinary thinking, is reliably available, and plays the sort of role an internal cognitive process would normally play, then merely pointing out that it is outside the skull may not settle whether the resulting capacity belongs to the person.

That is exactly the pressure on Sophie's case. "She cannot do it without AI" may show that her biological brain lacks a certain unaided skill. It does *not yet* show that the larger Sophie-plus-tool system lacks a genuine competence. The harder question is what kind of coupling would make the external resource part of her cognitive agency rather than merely a service she happens to hire.

Watch for four ideas

Parity

If the same functional role would count as cognitive inside the head, why should location alone disqualify it outside?

Coupling

How regularly, directly, and reliably must a person interact with an external resource before it becomes part of a standing capacity?

Otto vs. Inga

The notebook case contrasts biological memory with a trusted external store that guides action in a similar way.

Constitution vs. assistance

Is the tool merely helping a mind that remains internally complete, or partly constituting the cognitive process itself?



Open the assigned video

<https://www.youtube.com/watch?v=vrsQFSEtEqA>

Note: The guide links to the user-selected video but does not reproduce its transcript or claim a title/creator that could not be independently verified here.

After watching, answer one question: Is a typical LLM more like Otto's notebook, more like a collaborator, or neither? Name the feature that matters most: reliability, availability, automatic endorsement, generativity, verification burden, or something else.

Use in the packet: Link only. No video transcript or still is reproduced. The philosophical bridge on this page is BoMinds' summary of the extended-mind problem; the full Clark-Chalmers paper appears later as Material 4.

PRIMARY
MATERIAL

1/4

Aristotle, *Nicomachean Ethics*, Book II

Classic anchor · habituation, practice, and the formation of excellence · public domain U.S.

Aristotle is not writing about essays or AI. He is useful because he makes a simple but demanding claim: important human excellences are formed by repeatedly doing the kinds of actions in which those excellences are expressed.

The argument to understand

Book II argues that excellence is formed through practice rather than merely possessed as a dormant capacity. For this seminar, the most revealing part is Aristotle's contrast between producing a correct-looking result and possessing the knowledge or stable disposition that properly explains the result. That distinction maps surprisingly well onto AI-assisted work: a good output may be produced by chance or 'the suggestion of another' without showing that the learner owns the relevant competence.

VERBATIM SELECTED READING · BOOK II

Three complete selected passages from the public-domain Project Gutenberg translation, focusing on learning by doing, the quality of practice, and output produced by another's suggestion.

But the Virtues we get by first performing single acts of working, which, again, is the case of other things, as the arts for instance; for what we have to make when we have learned how, these we learn how to make by making: men come to be builders, for instance, by building; harp-players, by playing on the harp: exactly so, by doing just actions we come to be just; by doing the actions of self-mastery we come to be perfected in self-mastery; and by doing brave actions brave.

Again, every Virtue is either produced or destroyed from and by the very same circumstances: art too in like manner; I mean it is by playing the harp that both the good and the bad harp-players are formed: and similarly builders and all the rest; by building well men will become good builders; by doing it badly bad ones: in fact, if this had not been so, there would have been no need of instructors, but all men would have been at once good or bad in their several arts without them.

Now I can conceive a person perplexed as to the meaning of our statement, that men must do just actions to become just, and those of self-mastery to acquire the habit of self-mastery; "for," he would say, "if men are doing the actions they have the respective virtues already, just as men are grammarians or musicians when they do the actions of either art." May we not reply by saying that it is not so even in the case of the arts referred to: because a man may produce something grammatical either by chance or the suggestion of another; but then only will he be a grammarian when he not only produces something grammatical but does so grammarian-wise, i.e. in virtue of the grammatical knowledge he himself possesses.

Do not treat the analogy as identity. Aristotle is discussing arts and moral virtue, not generative AI. The useful question is whether some intellectual abilities, like his excellences, can only become ours through repeated first-person exercise.

Use Aristotle to ask

- Which parts of writing are merely production, and which are capacity-forming practice?
- Does a student need unaided practice before responsible delegation?
- Can AI itself become part of the practice rather than a substitute for it?

Important caveat

Aristotle's argument concerns moral virtue. Applying it to intellectual skill formation is an analogy supplied by this seminar, not Aristotle's own argument about writing.

Rights / use: The source work and this English translation are listed by Project Gutenberg as public domain in the United States. Source: Project Gutenberg eBook #8438. The selected passages are reproduced verbatim for educational discussion; public-domain status may differ elsewhere.

Project Gutenberg #8438 · [HTML full text](#)

PRIMARY
MATERIAL

2/4 A

Timothy Aylsworth & Clinton Castro, “Should I Use ChatGPT to Write My Papers?” (2024)

Philosophy & Technology 37:117 · Sections 3.1-3.3 · CC BY-NC-ND 4.0

Read the positive argument, not the abstract. Aylsworth and Castro's central claim is that some writing matters because it exercises capacities that partly constitute personal autonomy - our ability to examine reasons, choose ends, and govern our own lives.

Argumentative reading map · Section 3

1 · WHY AUTONOMY HAS SPECIAL VALUE

Section 3.1: The authors use a Kantian conception of “humanity” as the rational capacity to set and pursue one's own ends. Their point is not merely that autonomy is useful for getting other things we want. Rather, rational agents necessarily rely on this capacity whenever they ask what is worth wanting, compare reasons, or defend a view. That makes autonomy a candidate for a value that is not simply replaceable by efficiency, pleasure, or a better-looking product.

For Sophie's essay, this changes the question. If writing is one of the practices through which she learns to decide which reasons deserve endorsement, then outsourcing parts of that practice may matter even when the final sentences are excellent.

2 · DUTIES TO ONESELF DO NOT COVER EVERY SKILL

Section 3.2: The paper then narrows the claim. We do not have a duty to preserve every ability humans once performed unaided. Navigation, spelling correction, memory aids, and other tasks can often be outsourced without threatening self-government. The morally important distinction is between peripheral skills and capacities central to reflecting on ends, evaluating reasons, and revising one's conception of a good life.

This blocks an easy anti-technology argument. The authors do *not* claim that dependence on any machine is degrading. They claim that dependence becomes morally worrying when it prevents us from cultivating capacities that are constitutive of autonomy.

Short source phrase: “a duty to foster and safeguard your autonomy.” The rest of this page is BoMinds' argument-focused summary of Sections 3.1-3.2.

Use against Aristotle

Aristotle asks whether repeated first-person exercise forms the relevant capacity. Aylsworth and Castro add a normative claim: some of those capacities deserve cultivation because they help constitute self-government.

Pressure test

If Sophie delegates prose but still chooses her thesis, weighs objections, and revises her values, which autonomy-forming activity has actually disappeared?

Rights / access: The article is open access under CC BY-NC-ND 4.0. This page summarizes the argumentative core rather than reproducing the abstract or introduction. For the full source, read Sections 3.1-3.2 on Springer Nature.

DOI / open article · CC BY-NC-ND 4.0

PRIMARY
MATERIAL

2/4 B

Aylsworth & Castro: Why writing can be a special opportunity

Sections 3.3 and 4 · the positive argument, strongest objection, and reply

The paper's most useful argumentative work comes after it establishes the value of autonomy: it must still explain why writing one's own humanities paper is a sufficiently important opportunity to cultivate that autonomy - and why AI might not sometimes enhance it instead.

3 · THE "UNIQUE OPPORTUNITY" STEP · SECTION 3.3

The authors treat the duty to cultivate autonomy as an imperfect duty: it allows latitude, so missing one exercise is not automatically wrong. Their next step is opportunity-based. A student has special reason to do the work herself when a humanities assignment gives a good and unusually valuable opportunity to practice weighing reasons, scrutinizing inherited commitments, and receiving expert feedback on that reasoning.

On this view, a paper is not valuable because sentence production is sacred. It is valuable when the assignment forces the student to confront difficult reasons and make judgments she must own. If the task is mere boilerplate or administrative prose, the argument largely disappears.

4 · STRONGEST OBJECTION: AI CAN ENHANCE AUTONOMY · SECTION 4

The paper openly grants that technology can expand agency. AI can reduce tedious work, help generate objections, expose a student to more information, and improve the quality of thought. So the authors cannot simply infer "AI assistance" from "loss of autonomy." Their reply is developmental: there is a morally relevant difference between an expert using AI to extend a skill already mastered and a learner using AI in a way that prevents the underlying capacity from forming.

That reply makes the student's stage of learning central. It also creates a difficult counter-question for this seminar: if Sophie uses AI precisely to generate more objections, check her reasoning, and reflect more deeply than she otherwise could, then perhaps the tool is supporting - rather than replacing - the autonomy-forming part of the assignment.

5 · THE ARGUMENT IS ASSIGNMENT-RELATIVE

Aylsworth and Castro also place obligations on educators. If a writing task does not genuinely require students to wrestle with reasons, reflect on commitments, and benefit from serious feedback, then the moral case for insisting on unaided writing weakens. The paper is therefore best read as a claim about worthwhile educational practices, not as a blanket rule that all AI-assisted writing is morally impermissible.

Seminar hinge: The strongest version of their argument is not "AI makes you lazy." It is: *some difficult intellectual practices are part of becoming the kind of agent who can choose and revise ends for herself.* Ask whether Sophie's actual workflow preserves that practice.

Read in the original: Sections 3.3 ("Writing and the cultivation of humanity") and 4 ("Objections and Replies"). Open access under CC BY-NC-ND 4.0.

[Full article](#)

PRIMARY
MATERIAL

3/4 A

Tiegue Vieira Rodrigues, “The Apt Curation Model” (2026)

Philosophy & Technology 39:39 · Section 4 · CC BY 4.0 · core model

Rodrigues pushes in the opposite direction. If AI performs the sentence-level production, human authorship may survive when the human is the competent curator whose skill explains why the final text succeeds.

The argument to understand

Rodrigues offers the strongest counterweight in this packet. Instead of locating authorship in sentence production, the Apt Curation Model relocates epistemic agency to higher-order control: designing the inquiry, iteratively directing the model, evaluating outputs, and integrating them into a responsible whole. On this account, AI assistance need not eliminate achievement; what matters is whether success is attributable to the human curator's competence rather than luck or blind deference.

VERBATIM ARGUMENTATIVE SELECTION · SECTION 4

Three argument-bearing paragraphs retained from the prior packet: the shift from composition to curation, the authorship criterion, and the role of verification.

The dilemma outlined in the previous section highlights the inadequacy of any model that locates the author's competence solely in the act of composition. Such models inevitably fail because distant writing delegates precisely the low-level control of sentence formation to AI. The solution is not to abandon human agency, but to redistribute it. The primary epistemic work of the distant writer is not compositional, but curatorial.

We propose the Apt Curation Model, which views authorship in the age of AI as the proper steering of a generative process. According to this account, a text produced through distant writing is a genuine epistemic achievement worthy of the human author if its success is due to the author's adroit exercise of a set of higher-order curatorial virtues.

Furthermore, the charge of deference-induced overconfidence actually reinforces the importance of curatorial virtues rather than undermining them. The writer who falls victim to these bias lacks sufficient evaluative and architectural virtue in their work. The competent curator establishes critical distance through verification practices and by verifying information against established sources and testing arguments against each other.

Use this as the rival principle. Compare Aristotle's demand that the agent possess the relevant knowledge with Rodrigues' claim that the relevant competence may move upward into curation. Can both be true for different kinds of task?

Four curatorial virtues

- Architectural design
- Dialogical refinement
- Evaluative verification
- Integrative synthesis

The decisive test

Would the same human competence reliably produce a good result across new topics and imperfect AI outputs, or was this particular success mostly luck?

Rights / use: © The Author(s) 2026. Open access under CC BY 4.0, permitting sharing, adaptation, distribution, and reproduction with attribution and indication of changes. The source paragraphs above are reproduced verbatim; BoMinds commentary is separate.

[DOI](#) / [open article](#)

PRIMARY
MATERIAL

3/4 B

Rodrigues: What “apt curation” actually requires

Sections 4-6 · four virtues, the worked case, and the opacity objection

Do not reduce the Apt Curation Model to “prompting well.” Rodrigues makes authorship depend on a demanding bundle of higher-order competences that must explain why the final text succeeds rather than leaving success to luck.

1 · Architectural virtue

The human designs the intellectual project: formulates the research question, thesis, argumentative structure, dependencies, and likely objections. The core idea is that authorial control can reside in designing the problem-space even when sentence-level realization is delegated.

2 · Dialogical virtue

The human does not issue one command and accept the result. Competence appears in iterative interaction: diagnosing weaknesses, changing constraints, directing revisions, and learning how a particular model behaves across contexts.

3 · Evaluative virtue

This is the epistemically decisive virtue. AI output is treated as unverified intellectual material. The curator must detect errors, trace sources, test arguments, filter hallucinations, and possess enough domain knowledge to know when a polished answer is wrong.

4 · Integrative virtue

The human must turn heterogeneous outputs into a coherent whole: connect claims, preserve argumentative direction, integrate personal insight, and create a unified intellectual product rather than a collage of model completions.

The worked case: success through competence, not luck

Rodrigues' Dr. Evans example matters because it shows the four virtues working together. The researcher begins with a generic AI synthesis, narrows the inquiry, notices a promising connection, asks the model to develop it, checks the resulting references, detects a fabricated citation, replaces it with real scholarship, and integrates the revised argument into the larger project. The epistemic success is attributed to the researcher's ability to steer, test, and synthesize - not to the mere fact that a fluent answer appeared.

The opacity objection

A natural objection is that an opaque and probabilistic model breaks the link between the human's skill and the output: perhaps even careful prompting only produces lucky success. Rodrigues replies by shifting the standard from perfect prediction to cultivated know-how. Skilled users can learn which specifications, verification routines, and iterative strategies reliably improve outputs even when they cannot explain the model's internal computation. The real test is whether the user's competence reliably discriminates good from bad output and makes successful results more likely across new cases.

Scope matters

The model is strongest for *epistemic authorship* - philosophy, science, policy analysis, and other writing where truth, justification, and knowledge claims matter. Rodrigues explicitly leaves room for different standards in poetry and experimental literature, where architectural or aesthetic control may matter more than factual verification.

Apply this to Sophie: If she can set the inquiry, interrogate the model, identify bad arguments, verify sources, and integrate the result across unfamiliar topics, Rodrigues has a principled basis for calling the essay her epistemic achievement. If she merely approves polished output she cannot independently evaluate, the model does not rescue her authorship.

Rights / access: © The Author(s) 2026. CC BY 4.0. This page adds BoMinds summaries of Sections 4-6; the preceding page retains the packet's verbatim argumentative excerpts.

DOI / open article · CC BY 4.0

PRIMARY
MATERIAL

4/4

Andy Clark & David Chalmers, “The Extended Mind” (1998)

Analysis 58(1), 7-19 · classic philosophy of mind · summary only

Return now to the argument behind Material 0. Clark and Chalmers challenge the assumption that a genuine cognitive ability must be located entirely inside the skull, making ‘Sophie cannot do it without AI’ less decisive than it first sounds.

The argument to understand

The paper defends *active externalism*: under the right conditions, parts of the environment can play a constitutive role in cognition rather than merely supplying input. Their guiding parity idea is functional. If an external process plays the same kind of active role that we would unhesitatingly count as cognitive if it happened inside the head, its location outside the body should not by itself disqualify it.

The famous Otto-Inga case makes the point through memory. Inga remembers a museum address biologically. Otto, who has a memory impairment, carries a notebook that he reliably and routinely consults. The notebook is tightly integrated into his ordinary action in a way that resembles Inga’s stored memory, so Clark and Chalmers argue that the relevant cognitive system can extend beyond Otto’s biological boundaries. This does *not* immediately prove that an LLM becomes part of Sophie’s mind. A generative AI can invent new content, vary unpredictably, and demand verification in ways Otto’s trusted notebook does not. But the paper blocks one easy argument: “If she needs an external tool, the ability cannot be hers.” The real question becomes how the tool is coupled to her reasoning and what competence remains in the larger human-tool system.

NO VERBATIM EXCERPT REPRODUCED

An author-hosted draft is available online, but no explicit reuse license for the journal text was verified for this guide. We therefore summarize rather than reproduce.

Suggested reading: Sections 2-4, especially the move from ordinary external tools to the Otto notebook case. Ask which features of Otto’s relationship to the notebook would fail for a typical LLM.

Possible similarities to AI

- Constant availability
- Integrated into task performance
- External resource changes what the person can do

Possible differences from AI

- Generative rather than stored content
- Opacity and hallucination risk
- May require active distrust rather than automatic endorsement

Rights / use: No explicit open reproduction license was verified for the 1998 journal version. Summary only; no article passage is reproduced.

Journal DOI · Author-hosted draft

03 Stress tests and live-session plan

Change one fact at a time

Change one fact	What it tests
AI only fixes grammar and spelling.	Is authorship about sentence production or intellectual contribution?
AI creates the outline; Sophie writes every sentence.	How much does argumentative architecture matter?
AI drafts everything; Sophie verifies every claim and can defend it.	Can understanding and answerability replace direct production?
Sophie cannot explain one central paragraph.	Does a local failure of understanding break authorship of the whole?
A human tutor gives exactly the same suggestions as the AI.	Which differences between human help and AI help are morally relevant?
AI becomes perfectly reliable and permanently available.	If the tool never disappears, why must the ability remain inside the individual?

55-minute flow

Time	Stage	What happens
0:00-0:07	Case + first poll	Vote on whether the essay is Sophie's and state one criterion.
0:07-0:17	Small-group rule	Choose labor, control, understanding, formation, or responsibility.
0:17-0:28	Practice + autonomy	Use Aristotle and Aylsworth/Castro to ask what capacities are formed through first-person practice.
0:28-0:40	Curation	Use both Rodrigues pages to test whether architectural, dialogical, evaluative, and integrative competence can ground authorship.
0:40-0:50	Extended ability	Return to Material 0 and Clark/Chalmers: does the relevant ability have to be located inside Sophie?
0:50-0:55	Final position	Define the minimum human contribution required for "my work."

Prepare before joining

- Read the case, watch Material 0, and work through all four text materials. Focus on the argumentative sections flagged in the packet.
- Choose one primary criterion for authorship or achievement.
- Bring one example of a tool-assisted achievement you still consider fully yours.

Closing instruction: Your final position must say both (1) what makes the work Sophie's and (2) what makes the relevant ability Sophie's. Those need not receive the same answer.

04 References, access, and reading-selection rule

This packet is designed so participants can move from the short guide to the underlying texts. The selections and summaries are organized around the argument actually needed for the seminar, not around article abstracts.

Reading-selection rule. For public-domain or licensed sources, the packet prioritizes passages that actually advance an argument, develop a distinction, answer an objection, or work through a case. Abstracts and introductions are used only for orientation, not as substitutes for the philosophical reading.

Full references and access

Full references and access. Video primer: <https://www.youtube.com/watch?v=vrsQFSEtEqA>. Aristotle, *Nicomachean Ethics*, Book II, Project Gutenberg #8438. Aylsworth, T. & Castro, C. (2024), *Philosophy & Technology* 37:117, DOI 10.1007/s13347-024-00809-w. Rodrigues, T. V. (2026), *Philosophy & Technology* 39:39, DOI 10.1007/s13347-026-01038-z. Clark, A. & Chalmers, D. (1998), *Analysis* 58(1), 7-19, DOI 10.1093/analysis/58.1.7.

Suggested preparation order

1. Read Sophie's case.
2. Watch Material 0.
3. Read Aristotle.
4. Read both Aylsworth/Castro pages.
5. Read both Rodrigues pages.
6. Finish with Clark/Chalmers.

What to bring

One rule for authorship, one rule for ability, and one fact that would make you change either rule.

Licensing note. The packet states the verified reuse status of each source on its own material page. Links are provided so participants can consult the complete text or video.