■ Research Article

Rebuilding Public Reason in the Age of AI: Rhetorical Citizenship and Critical Doxic Literacy for Higher Education

Mika HIETANEN, Th.D., M.A. — Associate Professor, Centre for Languages and Literature, University of Lund, Sweden

FUTUREd, Volume 1, Issue 1 (2026)

Pages: 31 – 42

ISSN: XXXX-XXXX (print)
ISSN 2760-8271 (Online)

Keywords:
argument literacy; rhetorical citizenship; epistemic crisis; algorithmic persuasion; public judgement

Corresponding author(s):

Mika Hietanen
mika.hietanen@sol.lu.se
https://orcid.org/0000-0001-5569-9590

Abstract: Democracies face a crisis of public reason grounded in a disintegration of shared epistemic ground and amplified by digital polarisation and AI-mediated disinformation. Current pedagogy, focused on deconstructive critique or fact-checking skills, fails to address how algorithmic systems engineer plausibility at scale. Higher education must cultivate discerning judgement alongside critique. Drawing on rhetorical studies, this essay proposes a dual-axis framework for civic pedagogy: Rhetorical Citizenship (ethical communicative agency, coalition-building, and accountable ethos) and Critical Doxic Literacy (mapping tacit belief systems and analysing engineered plausibility). This approach retools classical concepts – eikos (plausibility) and doxa (shared premises) – to navigate hyper-eikotic digital environments where humans and algorithms co-produce a psychological feeling of truth decoupled from empirical reality. Using comparative case studies – Japan's techno-animism, India's digital sovereignty, and Brazil's tiered liability – the essay demonstrates how these competencies operate across distinct governance and cultural contexts. It addresses cognitive, ethical, and social dimensions of learning in an AI-mediated public sphere, preparing graduates to map competing architectures of belief and exercise provisional judgement in conditions of manufactured uncertainty.

1. THE COLLAPSE OF PUBLIC REASON IN A HYPER-EIKOTIC AGE

Democratic societies worldwide are undergoing a foundational collapse of public reason. Today's digital sphere is fractured by polarisation, algorithmic radicalisation, and the erosion of shared epistemic foundations (Vaidhyanathan, 2018).

The contemporary digital ecosystem has led to what I term *a hyper-eikotic reality.* In the Aristotelian tradition, *eikos* (plausibility) facilitated deliberation under uncertainty by relying on what was generally accepted in the absence of definitive proof. In the digital age, however, socio-technical systems exploit this rhetorical logic. Algorithms, including Large Language Models (LLMs) and social media recommendation engines, generate a sense of truth that is increasingly disconnected from empirical reality.

These systems prioritise resonance – what matches users' emotional triggers and cultural identities – over factual accuracy, creating self-sealing epistemic enclaves where contradictions become unintelligible. The consequences are particularly evident in the proliferation of synthetic media. The industrialisation of persuasion enables bad-faith actors to deploy sophisticated, personalised narratives at an unprecedented scale (Tufekci, 2017).

As human intent and automated content now mix, the provenance of information becomes unverifiable, leading to an epistemic misalignment. Surface-level features, like formal complexity and an authoritative tone, elicit misplaced trust in the absence of an accountable author.

Although higher education is tasked with cultivating civic citizens (Colby et al., 2003), current pedagogy lags, focusing on deconstructive critique or isolated technically oriented skills like lateral reading that regularly fall short against the self-sealing disinformation.

The proposed framework is theoretically grounded and informed by research that documents deficits in students' ability to recognise, evaluate, and respond to arguments within civic contexts (Andrews et al., 2006; Rapanta et al., 2013; Breakstone et al., 2021; Wineburg & McGrew, 2019).

This response consists of two interlocking axes. The first, *Rhetorical Citizenship*, reconceptualises civic identity as a discursive practice involving the cultivation of ethical communicative agency, coalition-building across difference, and the reclamation of *ethos* as personal accountability rather than mere performed credibility.

The second axis, *Critical Doxic Literacy*, offers the cognitive toolset necessary for mapping tacit belief systems (*doxic cartography*) and analysing how plausibility is engineered (*eikotic analysis*). These analytical competencies enable *comparative judgement* – the ability to provisionally weigh incompatible belief architectures. This often requires *doxic translation* to render differing frameworks mutually intelligible without requiring consensus.

The two axes function as complementary orientations rather than sequential stages and are ideally enacted within the same task.[1]

Three international case studies – Japan's techno-animism, India's data sovereignty framework, and Brazil's platform liability ruling – illustrate how the framework works across distinct governance and cultural contexts. The essay concludes by proposing the university as a laboratory for civic reason.

Three concepts from classical rhetoric recur throughout what follows; each is retooled rather than merely borrowed for the digital context. See Table 1.

Table 1 - Core rhetorical concepts

Term	*Classical meaning*	*Distortion*
Eikos	Plausibility: reasoning from what an audience finds probable	Algorithmically manufactured felt truth, decoupled from empirical reality
Doxa	Shared premises: often unstated assumptions structuring a community's common sense	Incompatible common senses (*doxai*) that coexist without shared ground in the public sphere
Ethos	The credibility a communicator projects through character, standing, and argument	*Synthetic ethos*: credibility engineered as pattern, without an accountable subject

2. RHETORICAL DISRUPTION

2.1. The Rhetoric of the Machine

The emergence of generative AI introduces a profound rhetorical disruption: *logos without ethos* (Gunkel, 2025). Whereas classical ethos is rooted

[1] *Rhetorical Citizenship* draws on the Copenhagen School (Kock & Villadsen, 2012, 2015). *Critical Doxic Literacy* builds on the analysis of the contemporary epistemic crisis developed in the author's forthcoming monograph, *Navigating Uncertainty* (manuscript in preparation).

in the speaker's character and historical standing, an AI's output is what I term *eikotic mimicry* – the simulation of credibility through surface markers of authoritative discourse. These systems optimise for persuasive fluency – what Bender et al. (2021) call 'ersatz fluency' – rather than factual provenance or moral accountability.

From a literacy perspective, students encounter AI-generated text that exhibits four deceptive markers:

- *Authorless authority* – text projects confidence through formal register (nominalisation, passive voice: 'It is widely considered that …') while occupying a 'position from nowhere' (Kjeldsen, 2025, p. 59) – no character, no social standing, and no accountability.
- *Citation mimicry* – fabricated references that imitate scholarly form without verifiable sources (Walters & Wilder, 2023).
- *Performative complexity* – more formally complex syntax, while less lexically diverse (Terčon, 2024), creating an illusion of expertise without conceptual understanding.
- *Hallucination as invention* – plausible-sounding falsehoods optimised for narrative coherence rather than truth, what Bender et al. (2021) describe as 'stochastic token prediction'.

These markers converge on the core literacy challenge: students misrecognise probability-driven text generation as rational reasoning. As Bender et al. (2021) explain, LLMs are 'stochastic parrots' that assemble sequences with high statistical likelihood in context – inhabiting the forms of argument (syllogism, scholarly citation, authoritative claim) while potentially filling them with empty or false content. The danger is not that AI produces bad arguments, but that it produces convincing-looking arguments whose plausibility masks their detachment from truth and human accountability.

This matters for literacy pedagogy because rhetoric – in the Aristotelian sense used here – presupposes audiences capable of evaluating arguments and speakers accountable for their claims (Aristotle, 2007, 1.1.1). Synthetic ethos severs this link: students encounter texts that perform argumentative responsibility without possessing it. The pedagogical project advanced here rests on teaching students to recognise this severance.

2.2. Engineered Plausibility

The current crisis stems from the transformation of the relationship between *eikos* and *doxa*. As Pierre Bourdieu (1977, pp. 164–170) defined it, *doxa* is the 'universe of the undiscussed' – foundational, unstated premises that structure our perception of the world. In a healthy public sphere, *doxa* provides the shared ground for deliberation. In the digital age, we see a collision of *doxai* – fundamentally different worlds of common sense that make cross-partisan communication nearly impossible. These *doxai* are maintained and amplified by systemic incentives, including platform architectures that profit from moral outrage and affective polarisation.

When algorithmic systems engineer *eikos* at scale – manufacturing plausibility tailored to pre-existing beliefs – they do not merely introduce false claims into a shared *doxa*; they fracture *doxa* itself into irreconcilable epistemic enclaves. This is the hyper-eikotic condition: a proliferation of mutually unintelligible plausibility systems, each internally coherent but collectively fragmenting the shared ground required for democratic deliberation. Traditional models of argument analysis and media literacy, designed for a world in which citizens inhabit a common evaluative framework, cannot address a crisis in which the frameworks themselves have multiplied and diverged.

3. BEYOND CRITIQUE

3.1. Weaponised Critique

For decades, the *hermeneutics of suspicion* has dominated, especially in the humanities. While vital for revealing injustices, deconstructing bias, and interrogating power structures (Graff, 2003), traditional deconstructive critique has reached its limits in the digital context (Felski, 2015). Deconstruction alone risks producing a citizenry adept at dismantling the arguments of others but equipped with few tools for constructive coalition-

building or for finding common ground in the face of 'wicked problems' like climate change or AI governance.

When students are taught to view every text as a site of manipulation and every claim as a power play, they risk emulating the tactics of disinformation agents who have co-opted the language of critical thinking.

Disinformation agents now frequently use the tropes of deconstruction – urging audiences to 'do their own research' or 'question authority' – to sow doubt about legitimate expertise in public health and democratic institutions (McIntyre, 2018). This weaponised critique leads to epistemic nihilism, where the inability to distinguish between cynical scepticism and healthy critique results in a paralysis of judgement, as seen, for example, in students' reliance on superficial credibility markers (Breakstone et al., 2021).

Research in Civic Online Reasoning has made progress in teaching 'lateral reading' and source verification (Doss, 2024). However, verifying a source is no longer sufficient when the narrative itself creates a self-sealing logic that renders the source irrelevant to the believer. As Wineburg and McGrew (2019) emphasise, the abundance of information on the web requires not just checking a source, but actively investigating who is behind the information and what institutional interests shape the narrative arc.

3.2. The Cognitive Dimension

A further obstacle is cognitive rather than curricular. Students tend toward cognitive closure – a preference for definite answers over discomfort with ambiguity. This blunts the kind of uncertainty-tolerant inquiry the hyper-eikotic environment demands (Kruglanski & Webster, 1996). Research on civic reasoning adds a complementary finding: students who develop the capacity to step beyond the particulars of a situation toward broader moral perspectives show stronger deliberative competence and identity development; yet this capacity is precisely what content-coverage approaches and algorithmically personalised information streams tend to suppress (Immordino-Yang et al., 2024; Gotlieb et al., 2024; Sunstein, 2017). The exercises proposed in this framework are designed to work against this disposition by asking students to map the conditions under which claims *feel* correct, rather than simply to identify what is correct – training in provisional, defeasible judgement rather than answer retrieval.

Together, these obstacles – the co-optation of critique and the cognitive pull toward closure – point toward the same need: a constructive supplement to deconstruction, one that replaces the move from dismantling claims with the cultivation of public judgement. The following two axes are designed to provide it.

4. FIRST AXIS: RHETORICAL CITIZENSHIP

4.1. Citizenship as Discursive Agency

To counter the disintegration of public reason, higher education must shift from a pedagogy of deconstruction to a pedagogy of *Rhetorical Citizenship* (RC). Drawing on the Copenhagen School, this framework defines citizenship as a discursive practice enacted through public communication (Kock & Villadsen, 2012).

Citizenship is enacted through the way we talk, write, listen, and argue in public fora (Crick & Lockyer, 2010). This requires moving beyond technical proficiency in communication – such as clarity and delivery – toward an empowerment model that treats students as emerging rhetorical actors with the agency to shape their digital environments. This axis emphasises the capacity for constructive coalition-building, framing issues for diverse audiences, and collaborative world-building.

This reframing aligns with the emerging concept of Critical Digital Pedagogy, which argues that education must foster *digital agency* – the ability to not just use tools but to interrogate and reshape the digital spaces one inhabits (Morris & Stommel, 2018; boyd, 2014).

4.2. Rights and Responsibilities

Rhetorical Citizenship involves two complementary dimensions: rights and responsibilities. As citizens, we have the right to

participate in public deliberation, but we also have a discursive responsibility to others – a duty to represent the ideas of others fairly and to engage in ethical argumentation (Kock & Villadsen, 2015).

This framework is operationalised through the *calibration of affect*. In a hyper-eikotic age, where algorithms profit from rage, RC teaches students to discipline their emotions, turning destructive affective polarisation into productive affective solidarity. It validates *pathos* as a legitimate part of public reasoning – such as righteous anger at injustice – while distinguishing it from the manipulative emotional contagion of the digital sphere (Brady et al., 2021).

This axis reclaims *ethos* as a relational and ethical practice. Rather than viewing credibility as a commercial performance or a brand, students are taught that ethos is constructed through accountability and the willingness to stand behind one's claims. This counters the authorless authority described in Section 2.1: when no identifiable author stands behind AI-generated text, accountability disappears, and communicative appeals become impossible. By learning to identify the absence of accountable ethos in algorithmic outputs and to construct it explicitly in their own work, students develop a critical counterweight to synthetic discourse.

These dimensions are also assessable; each corresponds to observable communicative behaviours:

- *Ethical framing*: represents opposing views fairly; signals accountability.
- *Coalition-building*: frames issues for diverse audiences and seeks overlap without erasing difference.
- *Affect calibration*: converts reactive outrage into productive solidarity while naming manipulative contagion.

5. SECOND AXIS: CRITICAL DOXIC LITERACY

5.1. Definition and Scope

While RC provides the ethical framework, the second axis, *Critical Doxic Literacy* (CDL), provides the cognitive toolset: *Doxic Cartography* (mapping tacit premises), *Eikotic Analysis* (how plausibility is engineered), and *Comparative Judgement*. I define CDL as follows:

> The advanced competence to map, analyse, and adjudicate between competing architectures of plausibility that constitute our public sphere, with a specific focus on identifying the tacit assumptions (*endoxa*), rhetorical strategies (human and synthetic), and systemic incentives (platform, economic, and governance logics) that maintain specific worldviews.

Unlike standard critical thinking, which focuses on individual arguments, CDL focuses on the systems that make those arguments feel true. It is the literacy of understanding how doxa is shaped, amplified, and weaponised by systems (algorithms, corporations, political movements). Table 2 indicates how CDL complements existing literacies.

Table 2 - Foci and Competencies of Literacy Types

Literacy Type	*Primary Focus*	*Key Competencies*
Critical Thinking	Soundness of individual arguments	Identifying fallacies, evaluating evidence, logical reasoning (Facione, 1990)
Critical Media Literacy	Representation and power in media	Analysing how media constructs reality; questioning ownership and representation (Kellner & Share, 2007)
Civic Online Reasoning	Credibility of online information	Source verification and lateral reading (Wineburg & McGrew, 2019)
Metaliteracy	Production and social media environments	Metacognitive reflection, collaborative production, and sharing in participatory environments (Mackey & Jacobson, 2011)
Critical Doxic Literacy	Architectures of plausibility and systemic incentives	Doxic Cartography, Eikotic Analysis, and Comparative Judgement

5.2. Doxic Cartography

The first core competency of CDL is *Doxic Cartography* – the ability to map the network of shared beliefs (*endoxa*) that constitute a

community's worldview. In a hyper-eikotic landscape, communities operate with fundamentally different common senses. Doxic Cartography asks:

- What are the foundational, often unstated premises this community holds as true?
- How do these premises interact to create a self-sealing sphere of plausibility?
- What is the hierarchy of authority within this doxa? (e.g. scripture, empirical data, personal intuition).

This skill is essential for *listening* across differences. Persuasion and even basic intelligibility presuppose a map of the other's doxa. It shifts the student from asking 'Is this fact true?' to 'Within what system of belief does this statement function as a truth?'

Example. When analysing vaccine hesitancy, Doxic Cartography does not merely debunk false claims. It maps the underlying doxa of distrust in pharmaceutical profit motives and ideals of natural purity. With that map, engagement strategies can be designed that address root assumptions rather than surface claims. This approach also draws on our findings that hate speech and radicalisation often rely on proximate markers – codes and symbols signalling adherence to a given doxa without violating platform rules (Eddebo et al., 2024). By mapping such codes, students delineate the boundaries of a community's belief system.

5.3. Eikotic Analysis

The second competency, *Eikotic Analysis*, focuses on the construction of engineered plausibility, examining how actors manufacture the feeling of truth, and how platform incentives function as rhetorical devices. Ethos, including *synthetic ethos*, is here crucial.

Central to this competency is analysing *synthetic persuasive content* – algorithmically generated representations designed to trigger the recognition of truth without grounding in accountable human judgement. Students are therefore taught to treat algorithms as rhetorical actors with their own engineered *ethos* and *pathos*. This involves analysing *procedural persuasion* – the way interactions like infinite scroll or 'likes' structure the arguments that can be made and the emotions elicited. It also includes identifying *corporate ethos repair*, where entities use ethics-washing or strategic partnerships to perform virtue while maintaining extractive models (Metcalf et al., 2019).

5.4. Comparative Judgement

The goal of CDL is *Comparative Judgement* – weighing incommensurate frameworks. Students ask:

- Which framework offers a more coherent, inclusive, and empirically resilient account of the world?
- Which framework relies more heavily on manipulative pathos or exclusion?
- How should conflicts between scientific and political *doxai* be adjudicated in policy contexts?

The first two competencies – Doxic Cartography and Eikotic Analysis – are analytical tools. Comparative Judgement is where those tools are put to work: the practice of weighing incompatible belief architectures against one another while holding one's own position provisionally.

This often requires *doxic translation*: rendering incompatible worldviews mutually intelligible without requiring agreement. Following Mouffe (2005), this fosters agonistic pluralism – the capacity to recognise opposing positions as legitimate rather than merely irrational. As the Japan case below demonstrates, techno-animist premises need to be translated into a Western framework to render the incompatibility itself intelligible.

At stake in Comparative Judgement is the cultivation of a specific epistemic virtue: *defeasibility* – the capacity to hold judgements provisionally, to weigh incommensurate frameworks against one another, and to maintain the conditions of ongoing deliberation when certainty is unavailable. Unlike conventional literacy approaches that treat education as a path toward greater certainty, CDL proposes a different orientation: the mark of the educated citizen is not

the ability to arrive at correct answers, but the capacity to reason rigorously in their absence.

This orientation draws on a classical rhetorical ideal: the citizen-judge of the Athenian assembly whose trained judgement constituted the goal of public reasoning (Aristotle, 2007, 2.18.1).

The broader educational aspiration is not the student who has mastered a checklist of cognitive skills. Rather, it is the citizen-judge who can perceive the architecture of the epistemic environment they inhabit, map the belief systems within it, and exercise situated judgement. This requires awareness of the conditions that shape the arguments they encounter and of the cognitive habits they bring to them.

6. THREE DOXIC CONSTELLATIONS

The following cases engage three distinct doxa-governance constellations: techno-cultural (Japan), sovereignty-economic (India), and constitutional-legal (Brazil). The conceptual vocabulary is Greco-Roman in origin. However, the infrastructural challenges are operationally global: manufactured plausibility, synthetic ethos, platform-driven polarisation, and the erosion of shared epistemic ground.

6.1. Techno-Animism and Ritual Interaction

Japan's approach to AI presents a productive case for Doxic Translation, as its underlying belief architecture differs from the dominant Western framework. Rooted in a syncretic blend of Shinto animism and Buddhist philosophy, a *techno-animist doxa* recognises the potential for spirits (*kami*) to inhabit both organic and inorganic objects (Mykhaylychenko et al., 2025), challenging the Western Frankenstein Complex – the subconscious fear that artificial creations will inevitably destroy their creators (Allison, 2006).

The *endoxa* sustaining this *doxa* can be mapped concretely. Japanese techno-animism rests on four premises: non-human entities can possess or develop spirit; relationships and rituals constitute meaning rather than merely express it; harmony between humans and non-humans is a civic good; and technological progress can be spiritually inflected rather than threatening. These premises underpin Japan's Society 5.0 vision (Cabinet Office of Japan, 2025).

Applying Eikotic Analysis to Society 5.0 reveals a second layer: Wright (2024) demonstrates that Japan's official AI ethics discourse performs virtue – invoking human-centredness and harmony, while structurally insulating extractive digital models from regulatory challenge. This is the *corporate ethos repair* identified in Section 5.3: the engineering of a plausibility frame in which ethical aspiration and commercial interest are rendered indistinguishable. Society 5.0 thus functions as both an object of Doxic Cartography (what beliefs make it persuasive?) and of Eikotic Analysis (how does it manufacture its credibility?).

The 2025 AInimism project operationalises the techno-animist doxa through a three-stage ritual interaction design: a keyword 'awakens' the object for analysis of its features; the AI then speaks as the object, drawing on stored memories; a closing transformation stores the exchange and reframes the human–nonhuman relationship (Mykhaylychenko et al., 2025; Mah et al., 2020). This sequence is itself a pedagogical resource. Students might first reconstruct the endoxa that make such a ritual plausible within a Shinto-Buddhist frame, then engage briefly with the protocol to observe how the framing shifts their own stance toward nonhuman agents. The exercise concludes with Doxic Translation: students reframe the rationale for this design for a Western audience shaped by the Frankenstein Complex, making the Japanese position intelligible without requiring agreement. The progression from mapping to enactment to translation illustrates how the three CDL competencies can operate in sequence.

6.2. Digital Sovereignty and the 2025 Rules

India's digital imaginary, centred on digital sovereignty, illustrates a middle path between state- driven and market-driven infrastructures, challenging the extractive patterns of data colonialism (Couldry & Mejias, 2019). In this doxa, data is viewed as a national asset, akin to oil or minerals (The Economic Times, 2019).

The Digital Personal Data Protection (DPDP) Act

of 2023 and its 2025 Rules operationalise this via the Consent Managers – India-based, neutral entities that provide a single portal for individuals, the Data Principals, to manage, review, and withdraw consent across multiple platforms (Pimentel et al. 2025; The Attorneys, 2025). The rhetorical shift from Data Subject to Data Principal is itself an object of analysis: it functions as an enthymeme whose suppressed major premise is that individual rights are a form of resource ownership, supporting India's push for Digital Public Infrastructure as a democratic alternative to Big Tech (The Economic Times, 2019).

Students working on this case analyse the enthymematic reasoning in the term Data Principal: by framing individuals as sovereign data owners rather than mere consumers, the term carries an unstated premise about rights and agency. They then construct a doxic map contrasting individual privacy rights with collective economic autonomy. The Doxic Translation exercise traces a user journey: a Data Principal withdraws consent across services (social media, fintech, health tracking) via India's Consent Manager. Students evaluate how platforms built on continuous data collection respond when consent is easily revoked – connecting abstract analysis of platform credibility to concrete governance consequences.

6.3. Tiered Liability and Democratic Survival

In June 2025, Brazil's Supreme Court (STF) established a fault-based liability model centred on unequivocal knowledge (*ciência inequívoca*). This tiered system presumes liability for monetised or boosted content – platforms are responsible for harmful content they have profited from – and for content distributed via automated networks, where knowledge is presumed (Baker McKenzie, 2025). The ruling mandates a proactive duty to remove evidently illicit content such as threats to democracy or hate speech (Massaro & Reis, 2025). The Court prioritises the survival of the public sphere over absolute platform immunity.

The STF ruling offers a compact case in Comparative Judgement. Students might begin by reconstructing the court's epistemic standard – what 'unequivocal knowledge' means as a threshold, and which platform behaviours cross it. A second movement asks students to articulate the same ruling from two incompatible *doxai*, one libertarian and one institutionally pro-democratic, weighing a Responsibility-with-Freedom thesis against the risks of preventive self-censorship (Baker McKenzie, 2025).

The exercise demonstrates that the same documentary facts – the ruling's text – produce divergent eikotic conclusions depending on which belief architecture processes them. The goal of CDL is not to reach consensus. Rather, it is mapping the conditions under which consensus becomes structurally impossible.

7. THE UNIVERSITY AS A LABORATORY OF CIVIC REASON

When the two axes operate together, the most generative exercises ask students to move across both within a single encounter: diagnosing the plausibility structure of an AI-generated text (CDL) and then drafting a response that reclaims ethos by naming its source and standing behind its claims (RC). The audits that follow are designed with this integration in mind. CDL provides the analytical lens: students map the epistemic architecture, identify weaponised plausibility, and trace the belief structures that make disinformation feel true. RC provides the communicative ethics: students ask what an accountable, audience-aware response to that architecture would look like, and practise constructing one.

7.1. Hallucination and Algorithmic Audits

The integration of RC and CDL involves moving beyond standard media literacy (including lateral reading) toward more active, red-teaming approaches to digital information.

In an *Algorithmic Audit*, students create simulated profiles to document how the platform feed narrows their doxa over time – from nuance to outrage. This renders platform incentives visible and measurable.

In a *Hallucination Audit*, students act as red teamers for an LLM, provoking it to generate false content by asking about non-existent topics. They

then analyse the output for markers of synthetic ethos – confidence without provenance, passive constructions, etc. This helps students understand that LLMs engage in *eikotic mimicry* – assembling sequences of words that resemble rational argument but are actually stochastic predictions.

7.2. Eikotic Fortresses and Doxic Translation

The *Eikotic Fortress Analysis* is a methodology for mapping the network of shared beliefs that constitutes a community's epistemically shielded worldview. Unlike standard fact-checking, which focuses on the truth-value of isolated claims, this analysis asks why a narrative *feels* true to a specific audience.

Example: Climate change denial. A traditional exercise might focus on empirical errors. However, this misses the deeper rhetorical architecture. As documented by Breakstone et al. (2021), only 4% of students identified ExxonMobil as the chief sponsor of the 'CO2 Science' website, demonstrating the efficacy of superficial credibility markers. In an Eikotic Fortress Analysis seminar, students map the plausibility structure that allows a community to reject scientific consensus while maintaining internal logical consistency. They identify underlying endoxa, such as 'institutions are inherently corrupt', and analyse how this rhetoric employs performative conflict to solidify in-group loyalty while rendering external fact-checks irrelevant. The outcome shifts from spotting lies to understanding belief architectures, enabling intellectual empathy for civic engagement.

To cultivate RC, students also engage in Doxic Translation, in which they identify a point of friction between two communities and translate the concerns of one into the values of the other.

For example, can the student articulate the vaccine sceptic's fear of corporate profit in a way that resonates with the pro-science advocate's value of transparency? This exercise forces students to inhabit the doxa of the 'other', fostering the intellectual empathy required for true democratic deliberation.

8. WRITING ACROSS THE CURRICULUM (WAC)

WAC programmes are a natural vehicle for contextualising CDL across specialised disciplines, grounded in the principle that writing is not only a mode of thinking but a mode of civic reasoning (Bazerman et al., 2005). Table 3 illustrates how the framework's core competencies translate into discipline-specific assignments.

Table 3 - Examples for WAC assignments

Discipline	*Focus of CDL Integration*	*Practical Example Assignment*
STEM (*Sciences*)	Public Science Rhetoric	How to communicate climate-risk probability without ceding authority to bad-faith actors (Oreskes & Conway, 2012).
Humanities	Strategic Narrative Architecture Analysis	Map weaponised historical narratives in current political discourse, e.g. nationalistic memory policies.
Business	Ethos Repair and Ethics Washing	Deconstruct AI-principles documents to see how they perform virtue to forestall regulation (Metcalf et al., 2019).
Computer Science	Algorithmic Bias Audit	Investigate how training-data selection (e.g. Western-centric datasets) embeds specific doxai into supposedly neutral code.
Law	Liability and Duty of Care	Analyse the Brazil STF 2025 fault-based liability model alongside global standards.

9. CONCLUSION

The epistemic pressures shaping contemporary public life place renewed demands on higher education, while also creating opportunities for pedagogical innovation. The framework proposed here – *Rhetorical Citizenship* developed alongside *Critical Doxic Literacy* – responds to this moment. It offers a conceptual toolkit that helps students perceive how plausibility is manufactured, how belief systems organise themselves, and how democratic dialogue might still be sustained despite the erosion of shared ground.

The perspectives and exercises outlined here gesture toward a concrete educational aspiration: graduates who can map the belief architectures within which they operate, distinguish synthetic from accountable ethos, communicate across doxic difference without requiring agreement, and hold their own judgements provisionally while still being capable of rendering them. This is not a checklist of cognitive skills but a civic orientation capable of exercising public reason in conditions of genuine uncertainty.

Higher education has always faced the task of equipping graduates to reason within the epistemic conditions of their time. The present conditions are distinctive in one respect: the architectures of plausibility that students must learn to navigate are not emergent properties of culture or politics but engineered outputs of systems optimised for engagement.

This framework is a theoretical proposal, not a tested curriculum. Its empirical grounding lies in the convergent literature on student deficits in civic reasoning it draws on; whether the specific competencies it cultivates can be measured is one of the questions a forthcoming instrument – the Civic Argument Literacy Test – is designed to address (Hietanen & Svedholm-Häkkinen, submitted).

The traditions of democratic education centred on *Bildung*, agonistic pluralism, and reflective citizenship remain as relevant as ever (Løvlie et al., 2003; Biesta, 2006). They now require, in addition, that the university become a laboratory for making those engineered architectures visible – and for cultivating in students the rhetorical and epistemic resources to reason within them without being captured by them.

REFERENCES

Allison, A. (2006). Millennial monsters: Japanese toys and the global imagination. University of California Press.

Aristotle (2007). On Rhetoric: A Theory of Civic Discourse (G. A. Kennedy, trans.; 2nd ed.). Oxford University Press.

Andrews, R., Bilbro, R., Mitchell, S., Peake, K., Prior, P., Robinson, A., See, B. H., & Torgerson, C. (2006). Argumentative skills in first year undergraduates: A pilot study. Higher Education Academy.

The Attorneys. (2025, Oct. 14). India's data privacy revolution: A deep dive into the DPDP Act. https://bit.ly/46poUkC

Baker McKenzie. (2025, July 1). Brazil: The Supreme Court (STF) establishes that Article 19 of the Brazilian Internet Legal Framework is partially unconstitutional. InsightPlus. https://bit.ly/4ai8snp

Bazerman, C., Little, J., Bethel, L., Chavkin, T., Fouquette, D., & Garufis, J. (2005). Reference guide to writing across the curriculum. Parlor Press.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT'21), 610–623. https://doi.org/10.1145/3442188.3445922

Biesta, G. J. J. (2006). Beyond learning: Democratic education for a human future. Paradigm Publishers.

Bourdieu, P. (1977). Outline of a theory of practice (R. Nice, trans.). Cambridge University Press.

boyd, d. (2014). It's Complicated: The Social Lives of Networked Teens. Yale University Press.

Brady, W. J., McLoughlin, K., Doan, T. N., & Crockett, M. J. (2021). How social learning amplifies moral outrage expression in online social networks. Science Advances, 7(33), eabe5641. https://doi.org/10.1126/sciadv.abe5641

Breakstone, J., Smith, M., Wineburg, S., Rapaport, A., Carle, J., Garland, M., & Saavedra, A. (2021). Students' civic online reasoning: A national portrait. Educational Researcher, 50(8), 505–515. https://doi.org/10.3102/0013189X211017495

Cabinet Office of Japan. (2025). Moonshot Goal 1: Realization of a society in which human beings can be free from limitations of body, brain, space, and time by 2050. https://bit.ly/4qtKQ5t

Colby, A., Ehrlich, T., Beaumont, E., & Stephens, J. (2003). Educating Citizens: Preparing America's Undergraduates for Lives of Moral and Civic Responsibility. Jossey-Bass.

Couldry, N., & Mejias, U. A. (2019). The costs of connection: How data is colonizing human life. Stanford University Press.

Crick, B., & Lockyer, A. (2010). Active citizenship: What could it achieve and how? Edinburgh University Press.

Doss, G. (2024). The effects of digital inquiry group curriculum materials on student civic online reasoning skills, The Great Lakes Social Studies Journal, 4(2), 38–48. https://doi.org/10.25035/glssj.04.02.07

The Economic Times. (2019, Jan. 18). Mukesh Ambani urges PM to take steps against data colonisation. https://bit.ly/402cuvj

Eddebo, J., Hietanen, M., & Johansson, M. (2024). Automatic identification of hate speech: A case-study of alt-right YouTube videos. F1000Research, 13, 328, 1–24. https://doi.org/10.12688/f1000research.147107.1

Facione, P. A. (1990). Critical thinking: A statement of expert consensus for purposes of educational assessment and instruction (the Delphi report) (ERIC Document No. ED315423). American Philosophical Association.

Felski, R. (2015). The limits of critique. University of Chicago Press.

Gotlieb, R. J. M., Yang, X.-F., & Immordino Yang, M. H. (2024). Diverse adolescents' transcendent thinking predicts young adult psychosocial outcomes via brain network development. Scientific Reports. https://doi.org/10.1038/s41598-024-56800-0

Graff, G. (2003). Clueless in academe. Yale University Press.

Gunkel, D. J. (2025). Large Language Models: Logos without ethos. In: A. Hess & J. E. Kjeldsen (Eds.), Ethos, Technology, and AI in Contemporary Society: The character in the machine (pp. 231–247). Routledge.

Hietanen, M. (manuscript in preparation). Navigating uncertainty: Reclaiming the eikotic for public reasoning.

Hietanen, M., & Svedholm-Häkkinen, A. M. (submitted). The Civic Argument Literacy Test (CALT): Foundations and development.

Immordino-Yang, M. H., Kundrak, C., Knecht, D., & Matthews, J. (2024). Civic reasoning depends on transcendent thinking. Social and Emotional Learning: Research, Practice, and Policy, 4, 100067. https://doi.org/10.1016/j.sel.2024.100067

Kellner, D., & Share, J. (2007). Critical media literacy, democracy, and the reconstruction of education. In D. Macedo & S.R. Steinberg (Eds.), Media literacy: A reader (pp. 3–23). Peter Lang.

Kjeldsen, Jens E. (2025). Ethos in the machine: The rhetorical character of debate AI. In: A. Hess & J. E. Kjeldsen (Eds.), Ethos, Technology, and AI in Contemporary Society: The character in the machine (pp. 58–80). Routledge.

Kock, C., & Villadsen, L. S. (Eds.). (2012). Rhetorical citizenship and public deliberation. Penn State University Press.

Kock, C. & Villadsen, L. S. (2015). Introduction Rhetorical Citizenship as a Conceptual Frame: What We Talk About When We Talk About Rhetorical Citizenship. In: L. Villadsen & C. Kock (Eds.), Contemporary Rhetorical Citizenship (pp. 9–26). Leiden University Press. https://doi.org/10.24415/9789400601918-001

Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: 'Seizing' and 'freezing'. Psychological Review, 103(2), 263–283. https://doi.org/10.1037/0033-295X.103.2.263

Løvlie, L., Mortensen, K. P., & Nordenbo, S. E. (Eds.). (2003). Educating humanity: Bildung in postmodernity. Blackwell.

Mackey, T. P., & Jacobson, T. E. (2011). Reframing information literacy as a metaliteracy. College & Research Libraries, 72(1), 62–78. https://doi.org/10.5860/crl-76r1

Mah, K., Loke, L., & Hespanhol, L. (2020). Designing with ritual interaction: A novel approach to compassion cultivation through a Buddhist-inspired interactive artwork. In Proceedings of the Fourteenth International Conference on Tangible, Embedded, and Embodied Interaction (pp. 363–375). https://doi.org/10.1145/3374920.3374947

Massaro, H., & Reis, D. (2025, Nov. 3). From Shield to Scrutiny: Brazil's Supreme Court Redefines Platform Liability. Global Network Initiative. https://bit.ly/4a4I9lV

McIntyre, L. (2018). Post-Truth. MIT Press.

Metcalf, J., Moss, E., & Boyd, D. (2019). Owning ethics: Corporate logics, Silicon Valley, and the institutionalization of ethics. Social Research, 86(2), 449–476. https://dx.doi.org/10.1353/sor.2019.0022.

Morris, S. M., & Stommel, J. (2018). An urgency of teachers: The work of critical digital pedagogy. Hybrid Pedagogy Inc. https://urgencyofteachers.com

Mouffe, C. (2005). On the political. Routledge.

Mykhaylychenko, D., Thasin, M., Baradari, D., & Mhungu, C. (2025). A(I)nimism: Re-enchanting the World Through AI-Mediated Object Interaction. arXiv. https://doi.org/10.48550/arXiv.2509.25558

Oreskes, N., & Conway, E. M. (2012). Merchants of doubt. Bloomsbury Press.

Pimentel, A. C., Saunders, D. P., Hebson, P., & Ying, J. C. (2025, Dec. 8). India's new privacy law is here: What you need to know. McDermott Will & Schulte. https://bit.ly/45WHnVv

Rapanta, C., García-Mila, M., & Gilabert, S. (2013). What is meant by argumentation competence? An integrative review of methods of analysis and assessment in education. Review of Educational Research, 83(4), 483–520. https://doi.org/10.3102/0034654313487606

Sunstein, C. R. (2017). #Republic: Divided democracy in the age of social media. Princeton University Press.

Terčon, L. (2024). Linguistic Characteristics of AI-Generated Text: A Survey. arXiv. https://arxiv.org/pdf/2510.05136

Tufekci, Z. (2017). Twitter and Tear Gas: The Power and Fragility of Networked Protest. Yale University Press.

Vaidhyanathan, S. (2018). Antisocial media: How Facebook disconnects us and undermines democracy. Oxford University Press.

Walters, W. H., & Wilder, E. I. (2023). Fabrication and errors in the bibliographic citations generated by ChatGPT. Scientific Reports, 13, 14045. https://doi.org/10.1038/s41598-023-41032-5

Wineburg, S., & McGrew, S. (2019). Lateral reading and the nature of expertise: Reading Less and Learning More When Evaluating Digital Information. Teachers College Record, 121(11), 1–40. https://doi.org/10.1177/016146811912101102

Wright, J. (2024). The development of AI ethics in Japan: Ethics-washing Society 5.0? East Asian Science, Technology and Society: An International Journal, 18 (2), 117–134. https://doi.org/10.1080/18752160.2023.2275987