

DOI: <https://doi.org/10.63332/joph.v6i4.4197>

A Conceptual Framework for Pedagogical Tensions in Algorithmic Education Beyond Personalisation

T Rugube¹, DW Govender²

Abstract

Educational experiences are increasingly shaped by the algorithms that underly the content, assessment, and guidance in learning systems; however, their potential and limits in pedagogical practice are relatively unexplored. Despite being imagined as personalisation technologies, these systems introduce some fundamental tensions with the pedagogical principles. This paper, through using a systematic literature review process, proposes a framework that distinguishes four fundamental tensions: a) Algorithmic Optimisation vs. Pedagogical Goals; b) Standardisation vs. Differentiation; c) Efficiency vs. Deep Learning; and d) Individual pathways vs. Social learning. Analysing how these tensions manifest in the existing Intelligent Tutoring Systems, Learning Management Systems, and AI chatbots revealed the pedagogical trade-offs that favour certain educational values over others. The framework concludes with design principles for navigating these tensions through pedagogical transparency, human-algorithm complementarity, and bounded algorithmic scope.

Keywords: content algorithms, pedagogical tensions, learning design, algorithmic education, personalised learning, educational technology

Introduction

Algorithmic personalisation presents a compelling narrative in education. Sophisticated computer models promise to give every student the exactly correct information he or she requires at precisely the moment it will be meaningful to them (Zhang et al., 2024; Naseer et al., 2024). Intelligent Tutoring Systems (ITS) accurately diagnose what students know and adapt instruction accordingly. Learning Management Systems (LMS) utilise adaptive quizzing which alters its response based on student competencies (Adiguzel et al., 2024; Ersozlu et al., 2024), while AI chatbots respond to any question students may ask and therefore meeting their individual needs (Labadze et al., 2024; Neumann et al., 2025; Liu et al., 2025). Ultimately, it is evident that technology can finally, accomplish true personalisation for all. Yet, to achieve these outcomes, we find that we need to make the necessary choices. This paper's argument is that algorithmic, educational technologies generate inherent tensions with pedagogical principles. Content algorithms' optimisation logic usually conflict with the core tenets of educational design. These include constraining engagement metrics, rushing students through competencies, and reducing the time it takes for students to master the learning content.

From constructivist perspectives, knowledge is constructed through meaningful activities like problem-solving (Piaget, 1952), but typically algorithmic systems try to minimize difficulty. From sociocultural learning perspectives, collaborative meaning-making is seen as central to

¹ University of KawZulu-Natal, South Africa, trugube@gmail.com

² University of KawZulu-Natal, South Africa, govenderd50@ukzn.ac.za



cognition (Vygotsky, 1978; Lave and Wenger, 1991), but algorithmic personalisation by definition makes learning experiences individual. Critical pedagogy emphasizes the learner as agent and processes of consciousness-raising (Freire 1970), but the algorithmic way restricts choice in the name of optimisation. These are not problems that better design can fix, but they are tensions, where to strengthen one pedagogical value simultaneously weakens the other (Verghis et al., 2025; Richter, 2025).

The Problem of Hidden Value Choices in Neutral Technologies

What is at stake in the current discussions of algorithmic education is that they embody a disturbing tendency: on the one hand, we have technologies that promise personalisation or optimisation, yet these commitments are not made visible or suffer little scrutiny. An ITS that pushes learners beyond concepts they have “mastered” i.e., get the right answer to practice problems, embodies a behaviorist epistemology that understanding learning as simply correct performance (Skinner, 1958). An adaptive system preventing boredom may also compromise metacognitive abilities (Verghis et al., 2025). A chatbot that responds immediately assumes that explicit instruction is better than discovery learning. These are pedagogical choices, and yet, they remain as technical architecture. This results in what Gallagher & Breines (2022) call an “algorithmic hidden curriculum”, systems that steer learning with values never chosen by educational communities.

Existing Gaps in the Literature

Research on algorithmic education is reflected by three main categories, each with important limitations. First, algorithmic enhancement, focuses on enhancing technical performance (Khasawneh, 2024; Wu et al., 2024), but often assumes pedagogical goals as “goods” without reflecting on how the optimisation can thereby introduce pedagogical costs. Second, critical work investigates algorithmic bias, surveillance and equity. Noble (2018), Eubanks (2018), and Baker & Hawn (2021) primarily addresses what algorithms get wrong rather than exploring what are the structural tensions within algorithmic pedagogy itself.

What is under-theorised is why doing the technological implementation of such algorithms in pedagogic spaces always creates resistance (Hazzan-Bishara et al., 2025). Such frictions are due to structural incompatibilities between algorithmic reasoning and pedagogical practice. Frameworks are required to support educators, policymakers and designers to acknowledge these tensions, and explicitly navigate trade-offs.

Recent Frameworks and Their Limitations

Recent literature proposes several frameworks for AI in education. UNESCO's AI Competency Framework for Teachers (2024) centers on teacher and student AI literacy and ethical values, emphasizing human-centered approaches (human rights, agency, accountability). Its core tension is teacher autonomy and equity versus algorithmic control, focusing broadly on national curricula. By contrast, PTF's unique contribution is its focus on classroom-level design tensions—what educational value algorithms optimize and what they sacrifice—rather than curriculum standards or policy.

The DAIGE (Dynamic AI Governance in Education) framework (Sangwa et al., 2025) adopts a multi-level governance perspective, identifying tensions such as institutional control versus teacher autonomy and academic integrity versus innovation, linking macro-policy with micro-level practices. DAIGE's educational focus is on school and system-level policy in higher education. PTF is distinct in analyzing pedagogical trade-offs at the level of learner interaction with content algorithms, rather than governance.

An AI Implementation Framework for K-12 (ILO Group, 2024) offers a district-level roadmap

stressing balance between innovation (tailored learning, efficiency) and oversight (privacy, safety, human control). Its core tensions include district priorities versus classroom needs and tool flexibility versus standardized protocols, focusing on administrators and IT infrastructure. PTF complements this by examining how algorithmic tools shape learning experiences, making explicit the pedagogical trade-offs that district policies might leave implicit.

An AI Ethics Literacy Framework (Yang et al., 2025) defines core AI ethics dimensions (fairness, privacy, accountability) linked to student learning outcomes. Its key tension is between technological utility and ethical awareness, with focus on higher education students. PTF's unique contribution is not solely teaching ethical concepts to learners, but providing educators a lens to analyze how AI-enabled systems balance or conflict with pedagogical goals in practice.

Table 1 below summarizes these comparisons.

Framework	Core Tensions	Educational focus	PTF's Unique Contribution
UNESCO AI Competency Framework for Teachers (2024)	Teacher autonomy and equity vs. algorithmic control; human-centric vs. data-driven education	Teacher and student AI competencies (K–12 and HE)	Human-centered policy/ professional training; PTF focuses on classroom-level design tensions and trade-offs.
DAIGE: Dynamic AI Governance in Education (Sangwa et al., 2025)	Institutional control vs. teacher autonomy; innovation vs. accountability	School and system-level AI policies (K–12/HE)	Multi-level governance and ethics; PTF provides a complementary focus on micro-level pedagogical choices.
K–12 AI Implementation Framework (ILO Group, 2024)	Innovation and personalization vs. safety and oversight; district objectives vs. classroom flexibility	District/school administration and infrastructure	Technical/administrative guidance; PTF highlights pedagogical implications of such AI tools in the classroom.
AI Ethics Literacy Framework (Yang et al., 2025)	Individual utility vs. ethical principles (fairness, privacy, accountability)	Student learning/competence (primarily higher ed)	Student-centered AI ethics knowledge; PTF focuses on systems-level educational outcomes rather than just learner attitudes.

These examples illustrate that while existing frameworks address broad issues (teacher skills, policy, ethics education), the PTF uniquely targets how algorithmic decisions manifest as tensions within the instructional process itself.

The Pedagogical Tensions Framework

This paper presents the Pedagogical Tensions Framework (PTF), as a tool for analysing and navigating structural conflicts between algorithmic systems and pedagogical principles. The framework identifies four fundamental tensions:

Algorithmic Optimisation vs. Pedagogical Goals - The tension between computationally measurable objectives and more complex, less quantifiable educational values.

Standardisation vs. Differentiation - The challenge of algorithmic personalisation operating within standardized categories, pathways, and metrics.

Efficiency vs. Deep Learning - The conflict between computational optimisation for speed and the slow, iterative cognitive processes essential for deep understanding.

Individual Pathways vs. Collaborative Learning - The tension between algorithmic individualisation and the fundamentally social nature of knowledge construction.

For each tension, this paper offers theoretical underpinnings, identifies implementations present in current systems (ITS, LMS and chatbots), discusses pedagogical tradeoffs and suggests design principles through explicit calibration.

The three contributions of the PTF are as follows. Analytically, it provides a way of systematically assessing current systems. Pragmatically, it aids in the design of new systems. In theory, it shifts the framing of algorithmic education from technical (“what should we optimise?”) to a normative (“what values should drive optimisation and what are acceptable trade-offs?”).

This paper focus on content algorithms, the systems deciding what content a user sees, when and how they see it, based on information about the learner. This spans from recommendation algorithms, adaptive learning algorithms to generative AI for educational content. The framework does not cover algorithms for assessment, student-risk predictive analyses, or administrative AI, as these warrant separate analyses.

Structure of the Paper

Section 2 lays the theoretical foundations. Sections 3 to 6 examine the four tensions. These sections articulate theoretical bases and envelope, experiential equivalents, and practical implications. Design principles are detailed in Section 7.

Conceptual Framework

Tensions, Contradictions, and Trade-Offs

Educational technology discourse commonly constructs challenges as "problems" with right answers or compromises that can satisfy all stakeholders. However, many fundamental issues in higher education technology are better understood as tensions, structural conflicts requiring ongoing negotiation through value choices rather than one-time solutions. Following paradox theory, researchers define tensions (or paradoxes) as persistent, co-existing demands that cannot be fully reconciled (Schaap et al., 2024). This differs from contradictions, which imply mutually exclusive goals, and trade-offs, which typically involve compensatory compromises. The concept of "tension" serves as the primary analytic category because it emphasizes the ongoing, dynamic interdependence of opposing pedagogical values rather than problems with single solutions.

This reframing has significant implications for how educators and institutions approach educational technology. When issues are treated as tensions rather than problems, it becomes clear that strengthening one educational value (such as efficiency) inherently weakens another (such as deep understanding). This approach aligns with paradox and polarity management literature, which argues that embracing both poles simultaneously is crucial for long-term adaptability (Bohunovsky et al., 2023; Kemp & Scoffham, 2022).

For instance, algorithmic personalization requires standardised representations—concept hierarchies, skill taxonomies, performance metrics, and learner models—even as genuine

pedagogy demands acknowledging each learner's unique differences. These conflicting pressures intensify as algorithms improve, illustrating how one pedagogical aim (individualisation) can undermine another (standardisation) unless explicitly managed. The Pedagogical Tensions Framework treats such algorithmic-pedagogical conflicts as paradoxical tensions to navigate rather than as isolated problems to solve, grounding this approach in contemporary organizational and educational theory.

Algorithmic Neutrality

Algorithms are not neutral instruments (Verghis et al., 2025). To begin with, they operationalise educational construct that is defining “understanding” in terms of correct response, “engagement” in time on task, or “readiness” with a predictive model necessarily abstracts rich pedagogical constructs into computable surrogates. Second, they encode epistemological assumptions about knowledge and evidence of knowing. Third, they configure learner agency regarding what to learn, in what order, at what pace and through which assessments. Finally, they are driven by a certain educational agenda. Prioritizing the fastest route to certification inherently de-prioritizes other goals such as depth over speed, context over consistency, and immeasurable growth over measurable results (Khasawneh, 2024; Wu et al., 2024).

The PTF repositions such factors as values in tension with algorithmic logic requiring explicit consideration of trade-offs (Richter, 2025; Wu et al., 2024).

Learning Theory and Algorithmic Logic

The pedagogical tensions discussed in this article come as a result of core philosophical disjunctions between established learning theories and the logics of computational optimisation. According to the behaviorist perspective of learning (Skinner, 1958; Thorndike, 1913), learning is observable change in behavior as a result of exposure to the environment with reinforced behaviors becoming learned. This operationalises in educational software as simple systems counting correct responses and supplying reinforcement, adapting task difficulty on performance based on subject's response. This crude operationalisation makes content algorithms implicitly behaviorist even when designers follow constructivist principles.

Cognitivist learning theory posits that learners construct knowledge but emphasizes internal mental processes—schema building, managing cognitive load, and metacognition (Richter, 2025). Even if algorithms could incorporate spacing effects or scaffolding they would have to optimise using behavioral proxies rather than direct internal cognitive processes. This leads to a fundamental paradox: behavioral measures may not always be reliable indices of cognitive states. Constructivist learning theory treats learners as active builders of knowledge who construct meaning through investigation, experimentation (Richter, 2025). Such constructivist tendencies can certainly be embodied by algorithmic systems, but algorithms are forced to provide structure for inherently open-ended procedures. They have to determine in advance which explorations will be regarded as useful, which solutions will be considered correct and what outcomes should be evaluated. Open-endedness central to constructivism becomes constrained algorithmic mediation.

Sociocultural learning theory posits that learning is inherently social, situated in cultural activities and facilitated through participation in communities of practice (Richter, 2025). Algorithmically personalised learning paths may undermine the collaborative meaning-making at the heart of sociocultural views diluting the social dimension.

Critical pedagogy promotes learner agency and critical consciousness, a vision of education understood not as domestication but liberation (Richter, 2025). Algorithmic pathways that aim for pre-set learning goals could constrain the exploratory, querying stance central to critical

pedagogy, one which may serve to reproduce rather than counteract existing power relations.

Scope of the Framework

The PTF addresses content algorithms shaping what students learn including Intelligent Tutoring Systems (ITS), LMS adaptive features, AI tutoring chatbots, recommendation engines in MOOCs. The framework excludes assessment and proctoring AI, predictive modeling for early warning systems and administrative AI which warrant separate frameworks.

Interrelationships among Tensions

The four identified tensions do not operate independently but are interdependent and mutually reinforcing. Emphasizing one dimension can intensify conflicts in others. For example, algorithmic systems optimized for efficiency (fast mastery paths) produce shorter learning sequences, sacrificing depth for efficiency. Simultaneously, highly individualized pathways make students more solitary, reducing collaborative opportunities. Resolving Efficiency versus Deep Learning tension (by speeding progress) exacerbates Individual versus Collaborative tension (students work alone more quickly).

Similarly, pursuing Algorithmic Optimization versus Pedagogical Goals often relies on standardized measures (test scores, completion rates) at the expense of unmeasurable values. This standardization feeds into Standardization versus Differentiation tension: to be optimized, content must fit fixed taxonomies, limiting true differentiation. Conversely, efforts to differentiate (branching paths) complicate algorithmic planning and may slow the system, affecting efficiency.

These interrelationships form a network of "tensions of tensions." Standardization versus Differentiation and Individual versus Collaborative are closely linked. Group learning requires finding common ground and synchronizing paths, conflicting with highly individualized adaptive content. Strong differentiation for each learner may undermine peer collaboration, intensifying Individual versus Collaborative tension. Algorithmic Optimization versus Pedagogical Goals intersects with Efficiency versus Deep Learning. Optimizing for measurable pedagogical goals (quiz scores) usually drives efficiency but often at the cost of slower learning.

In sum, these tensions form an interactive structure: neither strictly hierarchical nor completely orthogonal. They entail trade-offs where adjusting one "dial" shifts others. Increasing personalization (to resolve Standardization versus Differentiation) may deepen Efficiency versus Deep Learning conflict by requiring more complex content generation. Understanding these interrelationships helps educators anticipate that interventions addressing one tension may ripple through others—sometimes alleviating, sometimes exacerbating them.

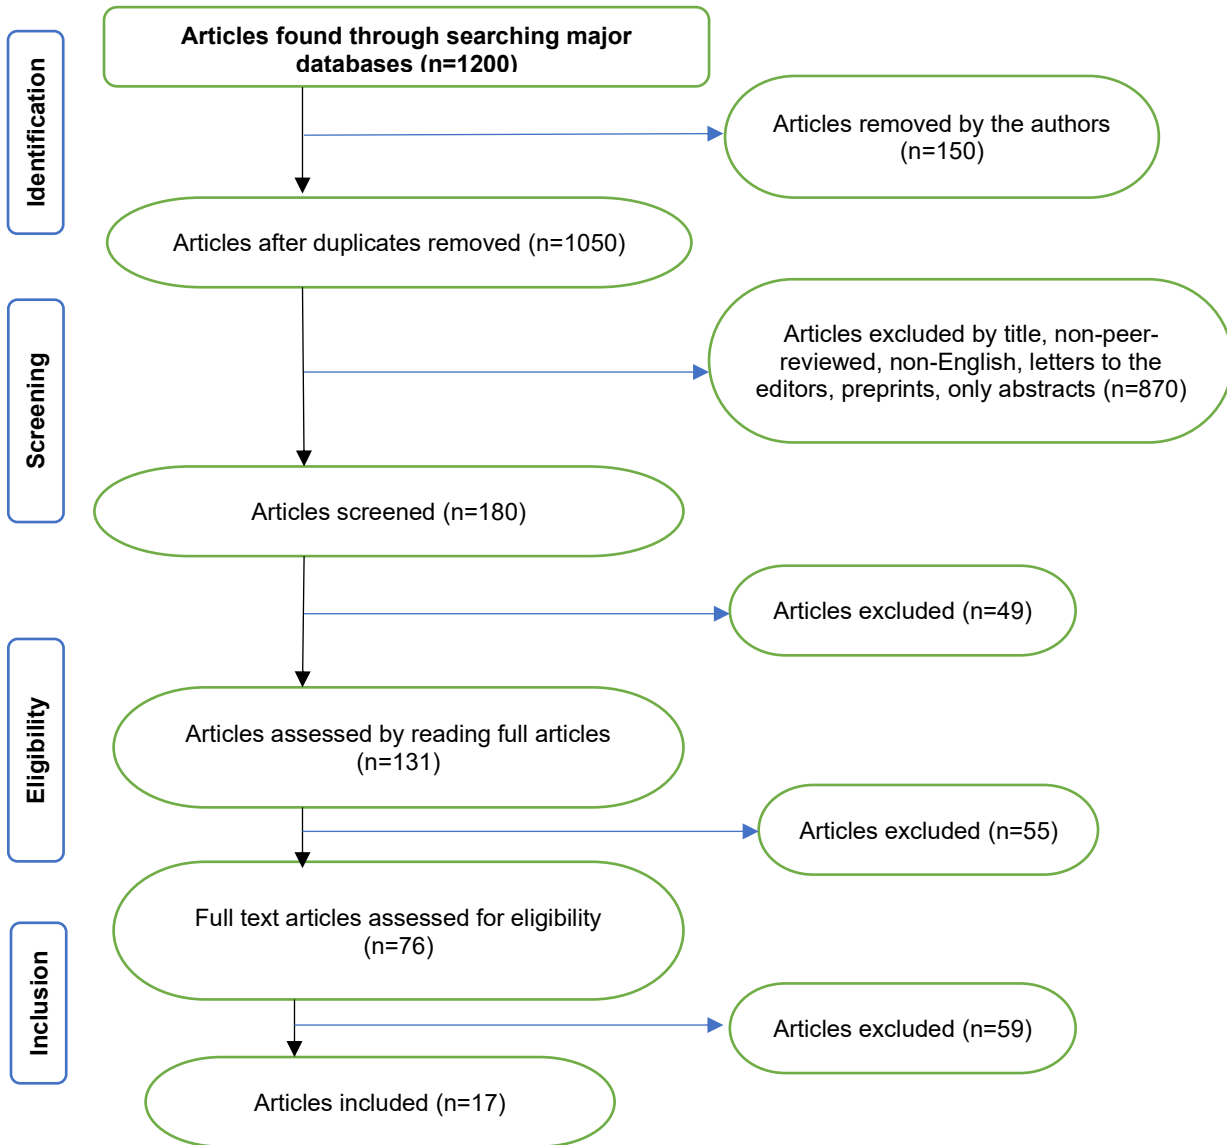
Literature Review Method

Literature Review Methodology

We followed the PRISMA 2020 guidelines for systematic reviews as adapted for educational technology research (Chan & Hu, 2023). Our search strategy was pre-specified and comprehensive. We queried major academic databases—Scopus, Web of Science, and ERIC—and key AI/education conference venues (e.g., AIED, EDM, CHI) using keywords such as algorithmic personalization, adaptive learning, pedagogical values, and educational AI ethics (Kasneci et al., 2023; Baker, 2022). We limited results to English-language, peer-reviewed journal articles and leading conference proceedings published between 2022 and 2024, reflecting current developments in AI-driven personalization and pedagogical tensions (Adiguzel et al., 2024; Demartini et al., 2024).

Eligible studies were required to examine algorithmic educational systems in higher education

and explicitly address pedagogical values or pedagogical trade-offs (Naseer et al., 2024; Verghis et al., 2025). Records focused solely on technical optimization, non-educational AI applications, or non-peer-reviewed commentaries—such as editorials or opinion pieces—were excluded to ensure conceptual and empirical relevance to pedagogical analysis (Richter, 2025; Holstein et al., 2023).



As outlined in recent methodological guidance on AI-in-education research syntheses, such as the reviews by Chan and Hu (2023) and Kasneci et al. (2023), the study selection followed the standard PRISMA phases (see Figure 1). During the Identification phase, the database searches yielded approximately 1,200 records. Consistent with large-scale searches reported in the AI

education literature—for example, Kasneci et al. (2023) noted similarly wide initial capture—about 150 duplicates were removed, leaving 1,050 unique records for screening.

In the Screening phase, two authors independently reviewed all titles and abstracts. Following practices described by Holstein et al. (2023) regarding rigorous review of adaptive learning research, we excluded studies that did not address algorithmic personalization in higher education, were not peer reviewed, or were otherwise out of scope (such as non-English or editorial formats). This process removed approximately 870 records, leaving 180 full-text articles for further assessment.

The Eligibility phase involved detailed full-text review. Echoing concerns raised by Richter (2025) and Verghis et al. (2025) about distinguishing pedagogically meaningful work from purely technical AI research, we excluded articles that focused solely on system optimization, performance metrics, or applications outside educational settings. After this assessment, 163 articles were excluded, resulting in 17 studies that met all inclusion criteria.

Finally, in the Included phase, two researchers independently coded the 17 studies. Drawing on procedures recommended in the educational data-mining literature—for instance, Baker (2022) emphasizes multi-coder verification in AI education reviews—we resolved all discrepancies through discussion to ensure reliability and conceptual coherence.

Throughout this process we maintained a transparent, predefined protocol (pre-specified search terms, databases, and inclusion/exclusion criteria) and documented reasons for exclusions at each stage. A PRISMA flow diagram summarizing the article selection process is provided in Table.

Author(s)	Year	Title	Focus/Contribution to Pedagogical Tensions Framework
Pelánek, R.	2024	<i>Adaptive Learning is Hard: Challenges, Nuances, and Trade-offs in Modeling</i>	Highlights that algorithmic adaptive learning models involve <i>inherent trade-offs</i> -- design choices create "numerous challenges, nuances, and trade-offs" -- urging designers to be aware that optimizing for one goal can negatively impact others (Pelánek, 2024).
Strielkowski, W., et al.	2024	<i>AI-driven adaptive learning for sustainable educational transformation</i>	Examines how AI-driven adaptive learning can make education more personalized, efficient, and aligned with sustainability goals. Emphasizes that personalized, accessible AI learning supports both individual growth and broader sustainable-education objectives (Strielkowski et al., 2024).
Demartini, C. G., et al.	2024	<i>Artificial intelligence bringing improvements to adaptive learning in education: A case study</i>	Describes an AI-supported learning analytics dashboard used in schools to identify at-risk students and tailor instruction. Shows how algorithmic tools can <i>augment teachers' decision-making</i> , enabling more equitable, need-based adaptations of learning processes

			(Demartini et al., 2024).
Naseer, F., et al.	2024	<i>Integrating deep learning techniques for personalized learning pathways in higher education</i>	Reports an experimental study (n≈300) where an AI-driven adaptive platform raised student grades and engagement by ~25% vs. traditional instruction. Demonstrates strong empirical gains from algorithmic personalization while also providing a scalable implementation framework (Naseer et al., 2024).
Neumann, A. T., et al.	2025	<i>An LLM-driven chatbot in higher education for databases and information systems</i>	Evaluates "MoodleBot," a course-specific GPT chatbot. Finds 88% accuracy on course queries and positive student acceptance. Concludes that LLM chatbots can <i>enhance personalized support</i> and reduce instructor workload in CS classes (Neumann et al., 2025), highlighting efficiency gains vs. pedagogical needs.
Sangwa, S., et al.	2025	<i>Dynamic AI governance in education (DAIGE): A multi-level framework</i>	Proposes and validates a four-pillar governance model for ethical AI in education. Shows that institutions using <i>inclusive, iterative policies</i> (multi-stakeholder co-creation) achieved 21% higher teacher/student buy-in and 50% fewer cheating incidents. Emphasizes that good governance can reconcile algorithmic innovation with pedagogical values (Sangwa et al., 2025).
Favero, L., Pérez-Ortiz, J. A., & Manrique, D.	2024	<i>Enhancing critical thinking in education by means of a Socratic chatbot</i>	Develops a Socratic-style chatbot (fine-tuned LLM) that prompts students to reflect on multiple perspectives rather than giving direct answers. In comparison tests, the Socratic tutor significantly boosted students' reflection and critical thinking versus a standard answer bot (Favero et al., 2024), countering the "efficiency" bias of answer-providing AIs.
Holstein, K., Alevan, V., & Rummel, N.	2023	<i>Co-designing adaptive learning technologies with teachers: Balancing personalisation and classroom</i>	Highlights that involving teachers in design helps balance personalization with classroom orchestration. By co-designing adaptive tools, educators ensure algorithms support learning goals without undermining group instruction, avoiding "personalization at the expense of shared

		<i>orchestration</i>	learning contexts" (Holstein et al., 2023).
Jansen, R. S., van Leeuwen, A., & Janssen, J.	2023	<i>Metacognitive scaffolding and self-regulated learning in digital environments: When support becomes dependency</i>	Shows that <i>excessive algorithmic scaffolding can undermine learner autonomy</i> . Their findings warn that continuous AI support may obstruct development of students' metacognitive and self-regulatory skills, creating dependency (Jansen et al., 2023). This reveals the trade-off between short-term efficiency (support) and long-term depth (self-regulation).
Kasneci, E., et al.	2023	<i>ChatGPT for good? On opportunities and challenges of large language models for education</i>	Provides a balanced overview of LLMs in education: Opportunities include generating content, boosting engagement, and personalization; Challenges demand new literacies (critical thinking, fact-checking) and raise bias and oversight concerns. Emphasizes that maximizing <i>efficiency and personalization</i> with LLMs must be weighed against maintaining critical, reflective learning and fairness (Kasneci et al., 2023).
Richter, F.	2025	<i>The ethical acceptability of personalisation via intelligent systems in education</i>	Critiques naive personalization: notes that, despite broad enthusiasm, many AI systems <i>"fail to capture the complexity of learning"</i> and neglect pedagogical context (Richter, 2025). Argues we must scrutinize personalization approaches ethically, ensuring algorithmic designs align with rich learning goals rather than treating education as a one-dimensional optimization.
Verghis, A. M., et al.	2025	<i>Beyond personalisation: Autonomy and agency in intelligent systems education</i>	Argues that typical personalization often undermines student autonomy (labeling it <i>"extrinsic personalization"</i>). Calls for a <i>reclaiming of personalization</i> around learner agency, complexity, and ethical self-direction: systems should invite curiosity and reflection, not just predict what students likely want (Verghis et al., 2025). This work reframes personalization as a balance between support and fostering independence.

Chan, C. K. Y., & Hu, W.	2023	<i>Students' voices on generative AI: Perceptions, benefits, and challenges in higher education</i>	Surveyed students on generative AI (e.g. ChatGPT). Found a positive attitude: students valued AI for personalized writing help and learning support. But they raised significant concerns about accuracy, privacy, and ethical issues, and worried AI use might hinder deep understanding. Thus, student perspectives highlight the tension between <i>AI's efficiency/personalization</i> and maintaining critical learning quality (Chan & Hu, 2023).
Peng, H., Ma, S., & Spector, J. M.	2019	<i>Personalized adaptive learning: An emerging pedagogical approach enabled by a smart learning environment</i>	Conceptualizes Personalized Adaptive Learning (PAL) as a new pedagogy: smart learning tech monitors learners in real time and automatically adjusts instruction to individual differences. Emphasizes that PAL represents <i>technology-empowered pedagogy</i> where adaptive algorithms tailor teaching strategies on the fly (Peng et al., 2019). This provides a foundational framework for understanding how algorithmic systems can serve pedagogical differentiation.
Baker, R. S.	2022	<i>Algorithmic bias in education</i>	Comprehensive review of bias in educational algorithms. Catalogs how different student groups (race, gender, SES, etc.) are affected at each development stage. By solidifying "which groups are impacted" and how, this work highlights a key pedagogical tension: even well-intentioned personalization can embed unfairness. It ends by proposing moving from mere fairness metrics toward equity (Baker, 2022).
Xu, X., et al.	2025	<i>AI optimization algorithms enhance higher education management and personalized teaching through empirical analysis</i>	Empirical study of AI optimization in higher education. Finds that AI-driven optimization " <i>effectively solve[s] complex educational problems and enable[s] personalized learning experiences</i> ", with ~25% gains in outcomes, engagement, and efficiency vs. traditional methods (Xu et al., 2025). Also notes serious challenges (data privacy, bias, need for human-AI

			interaction) in deploying such systems. Highlights trade-offs of powerful personalization: large benefits but new ethical/quality risks.
Nelson, A., & Elias, K. L.	2023	<i>Desirable difficulty: Theory and application of intentionally challenging learning</i>	Highlights Desirable Difficulties theory: intentionally adding effortful challenges (slower, "harder" tasks) can impede short-term performance but boost long-term learning. This concept underpins a key tension: not all algorithmic shortcuts are pedagogically beneficial. Difficulties that slow efficiency may actually foster deeper understanding (Nelson & Elias, 2023), suggesting algorithms should sometimes <i>withhold ease</i> to achieve lasting learning.

Tension 1: Algorithmic Optimisation vs. Pedagogical Goals

Theoretical Tension

Algorithmic optimization requires measurability. A useful objective function is said to have computationally maximizable learning goals (Zhou et al., 2025). Systems work towards quantifiable than measurable goals: test scores, completion rates time-on-task, problem-solving efficiency or engagement metrics (Baker, 2022). However, many of the most important educational aims are difficult to assess: critical thinking, creativity, wisdom, moral reasoning, aesthetic appreciation, identity formation (Bouckaert 2023; Thornhill-Miller et al 2023). Even when proxy measures are available, they only cover partly dimensions , such as the detection of critical thinking in written answers.

What is measurable can be managed. Wherever there are proxies for measurement, the algorithms drive those dimensions forward and where objectives are not measurable, they shrink , leading curriculum to focus on narrow skills that can be easily measured (Stanford 2023; Morgan 2025; FairTest, 2023).

Manifestations in Practice

LMS adaptive quizzing adapts the questions according to student answers. Canvas yields harder scaffolds following right answers and easier ones after wrong answers while keeping the mind in appropriate challenge zones (Strielkowski et al., 2024). This maximizes sustained involvement (as opposed to boredom and frustration) and efficient assessment. However, it forgoes desirable difficulty, how the tasks at hand are generating learning demands that will allow superior long-term retention but at a cost of retrieval decrement (Nelson & Elias, 2023; Pyke et al., 2025; Kinsey, 2023). It also reverses metacognitive calibration; students require practice in assessing their level of understanding, not systems that are continually adjusting the difficulty (Fan et al., 2025; Wang et al., 2025).

Generative AI platforms such as Khanmigo are built to mimic Socratic-style tutoring, which focuses on steering students towards solutions through the strategic use of questions, rather than just giving direct answers. To accomplish this, these systems consider answer relevance, accuracy and develop pedagogically moving hint sequence that would preserve student engagement (DiCerbo, 2025; Microsoft News Center, 2025). However, true Socratic pedagogy requires far more than this. Its nature demands a long form of conversation where the meaning

reveals itself over time, a process that requires patience well beyond what typical chatbot conversations can maintain (Fakour & Imani, 2025).

Furthermore, what is referred to here as Socratic teaching includes a type of uncertainty which implies that the teachers really work on issues with students (Favero et al., 2024). It also means pushing a student's thinking, something that may be momentarily frustrating but ultimately leads to lasting learning (Nelson & Eliaz, 2023). While contemporary AI text generation models are tuned to favor "useful," "factual," and entertaining responses, they must flatten these emphatically rich Socratic goals into no more than parody (Fakour & Imani, 2025; Kasneci et al., 2024).

Pedagogical Costs

The problem is that costs are compounded when systems focus on only measurable targets. "The curriculum shrinks to those things that can be tested quickly" as "teachers cut out any untested material" (Center for American Progress, 2023; Parcerisa & Pagès, 2025). Students develop an instrumental motivation towards extrinsic rewards rather than intrinsic interest (Jian & Liu, 2023). "The potentially infinite capacity of human beings to know is inappropriately squelched: immeasurable dimensions, deep rather than superficial understanding, aesthetic appreciation, moral reasoning and creative synthesis" (Vincent-Lancrin et al., 2019). Educators' professional judgement as to what constitutes meaningful learning is reduced to an algorithmic measure; data-driven prescriptive process superseding pedagogical discretion and thereby eroding teacher self-regulation (Hall et al., 2025). What becomes the most important question in education changes from asking what are valuable learning experiences to asking what can be measured and increased (Lane, 2020).

Navigating the Tension

Since quantifiable metrics certainly represent incomplete proxies (van Haastrecht et al., 2025), these tensions are mediated by several strategies:

Multi-Objective Optimisation

Rather than assessing educational AI systems according to one metric—whether students answer questions correctly or complete work rapidly, for example—we may want to consider them with respect to multiple dimensions that reflect what really counts in learning (Qian et al., 2025). For instance, we might attend to whether students are actually making sense of the material, whether they're building up a stamina for continuing on when there are hiccups (Kapur & Bielaczyc, 2012; Lynch et al., 2022), as well as if they're pushing themselves to reflect on their own thinking processes (Xu et al., 2025; Sumadyo et al., 2021) or try other strategies besides always relying on one trajectory. When we measure learning in this way, we can freely acknowledge the trade-offs at play, and accept, for example, that slowing down to let students wrestle productively with concepts may result in lower immediate test scores but deeper understanding over time. Lacking this kind of wider view, we expose ourselves to the possibility that some single goal, such as speed or accuracy, may slowly but surely capture and mold the whole system, despite our best intentions never imagined (Rahiman & Kodikal, 2025; Marcolino et al., 2025).

Human Designation of Unmeasurable Value

Education is not all that can be measured by algorithms, and this is where human judgment becomes indispensable' (UNESCO, 2024). Computers can't compute what teachers and students can protect. For example, they might reserve some "essential experiences" that all would regardless take part in irrespective of what the algorithm proposed. They might elect to carve out time for students to self-select work on topics or questions they are curious about, no guard rails provided by an AI system and create time for reflection times in which a student simply

thinks about learning without needing to produce evidence of it (Greene, 2021).

All this is consistent with the UNESCO AI Competency Framework (2024), which underpins that teachers need to have the authority to supersede algorithmic recommendations when their professional judgment indicates otherwise. Teachers will have a knowledge of their students that data thus far cannot, and by keeping this human oversight in place we make sure there is someone responsible for the educational experiences of these students, which contributes to trust in the system as a whole (Georgieva et al., 2025; Molenaar, 2022).

Algorithmic Humility

AI systems should be honest about their capabilities (Farrow, 2023). For example, a math tutoring chatbot might say to a student: “I can help you find an efficient method of solving this problem, but I cannot determine whether working through the problem on your own would allow you to appreciate the elegance of mathematics. ‘You are going to have to be honest with yourself about that.’ What this kind of honesty does, though, is to remind students not everything that matters about learning shows up in the data.

New methods in explainable AI may help to expose and render such limitations transparent (Gunasekara & Saarela, 2025) and subject-specific explanations led teachers to trust the technology more (Feldman-Maggor et al., 2024). Critically, teachers should serve as co-creators of such systems and not mere end users and AI is expected to reveal when its predictions are uncertain (Duan et al., 2024). There is also a fine line to tread: while transparency can be beneficial, overreliance on AI feedback may cause what researchers term “metacognitive laziness” concerning thinking critically about their own learning process having the AI do this for them (Fan et al., 2024).

Regular Metric Auditing

Educators should also reflect on what their AI systems incentivize and disincentivize that is, what kind of learners students would be if they optimally internalized the goals of the algorithm (Costanza-Chock et al., 2022). This is about what educational values the system reinforces and which ones it erodes. When it is clear there are misalignments between the algorithm's behaviour and real learning goals, these metrics should be adjusted (Mökander et al., 2022). Because algorithms engage in open-ended optimisation, continuous monitoring is needed to keep them advancing educationally meaningful rather than gaming the metrics (Baker & Hawn, 2021; Paige & Amrein-Beardsley, 2020).

Tension 2: Standardisation vs. Differentiation

Theoretical Tension

Educational AI systems confront a basic contradiction: In order to personalize learning, they must first standardize it. Algorithms, in the name of customizing education to every student, can only do so by sorting students into predetermined groups and tracks (Fisher, 2024; Zhu et al., 2021). Which is to say AI systems need standardized frameworks, such as learning style categories, and assessment rubrics to function effectively.

This tension between uniformity and individualisation plays out in two main dimensions. Well, adaptive learning systems generate paths to map students' selections and performance (Peng et al., 2019). However, all their choices need to be pre-designed and driven by "pre-determined characteristics" rather than being a response to emerging student needs (Van den Berg & Timmermans 2021). This results in “bounded adaptation”, systems adapt to student nuances only across the dimensions designers planned for (Dziuban et al., 2016). Unforeseen learning needs, interests or pathways are not supportable.

AI chatbots seem to bypass these restraints by producing what appears as an infinite number of

replies. However, GPT-based tutors generate responses by identifying patterns in training data, essentially learning what typically follows certain inputs (Lipenkova, 2023). They model overarching tendencies across many cases rather than generating novel explanations for each learner. These systems are developed to maximize the statistical likelihood, instead of pedagogical appropriateness for a student based on modeling how students were trained (Thrush et al., 2024). The personalisation is also bound as always by training data, not pre-programmed pathways, but it is bound all the same.

Pedagogical Costs

This tension between normalisation and variance generates a number of issues for learning. Knowledge is broken up: rich, connected experience gets torn apart at the seams into pockets of competency; interrelationships among concepts are severed (Horneber & Laumer, 2023). Second, students are reduced to data points that may be placed in pre-established categories of learning styles, and competence levels (Edwards & Cheok, 2018). Third, there may be a danger of cultural imperialism in the south driven by global standards, akin to politicization and cultural hegemony propagated by globalization (Selwyn, 2021). Systems that support only expected learning patterns also choke off serendipitous discoveries and self-directed exploration (Panagiotopoulos & Shahrezaye, 2024).

Navigating the Tension

There are several approaches to tackling this issue. First, teachers must separate what can be equitably standardized from what needs to remain unique. Fundamental concepts can have a common core of structure among all students, and application contexts, motivational factors, and knowledge linkages need to be individually tailored for each learner (Xie et al., 2019; Tlili et al., 2023).

Second, instead of creating pre-defined categories, it is also important to let learners identify their learning dimensions – tagging interests, experiences or goals on their terms. A dilatator may also be used in this bottom-up fashion to encourage acceptance of active agency while still allowing for pattern recognition (Kaur & Khamparia, 2022; Budhwar, 2024).

Thirdly, systems need to need to be transparent about their limitations. Where "personalized" pathways are provided they should foreground the idea that these are selections from pre-configured lists, not open-ended tailoring and thus enable learners to make informed choices about whether to go with recommendations or seek alternatives (Kumar et al., 2024).

Finally, platforms require 'escape valves' that allow learners to signal when algorithmic categories don't work, requesting alternative pathways, flagging mismatched content, or momentarily turning off recommendations altogether (López-Pernas et al., 2023; Papadakis et al., 2024).

Tension 3: Efficiency vs. Deep Learning

Theoretical Tension

Most algorithmic designs optimise for the path of least resistance, through training time to proficiency, avoiding frustration and speeding up skill development. Now those are noble aspirations, but deep learning typically involves something that may appear to be inefficient. Learners require time-consuming, yet unproductive struggle (Steenhof et al., 2023), temporary confusion leading to conceptual breakthroughs (Chowrira et al., 2023), desirable difficulties for better retention in the long-run (Nelson & Eliaz, 2023), and interleaving to prevent premature automaticity (Hwang et al., 2024).

The issue is this: algorithms that have been optimized for smooth, efficient learning can sometimes delete the friction that helps you remember better. Not all cognitive load interferes

with learning; some is the useful for learning (Greenberg & Zheng, in press). The problem is that productive struggle and unproductive struggling can be nearly indistinguishable to an algorithm since both take the form of poor performance in the short term.

Manifestations in Practice

Spaced repetition systems algorithms optimising review timing to minimize study time while maximizing retention. Cards appear just before forgetting, efficiently reinforcing memory. However, this prioritizes retention over initial encoding depth. Research shows generation effects, struggling to produce information rather than passively reviewing it—creates stronger memory traces even if initially less efficient (Kornell & Bjork, 2008). By optimising review efficiency, these systems may sacrifice encoding depth.

Intelligent hints in tutoring systems provide scaffolding when students struggle. Systems like Cognitive Tutor offer increasingly specific hints: from general problem-solving strategies to step-by-step solutions. This prevents extended frustration and keeps students progressing. However, Kapur's (2014) productive failure research demonstrates that struggling through problems before receiving instruction produces better learning than scaffolded support throughout. The hint systems, optimising to prevent struggle, may inadvertently prevent the productive failure enabling robust understanding.

Real-time feedback in LMS quizzes immediately indicates answer correctness. This enables rapid error correction, preventing students from practicing incorrect procedures. However, delayed feedback can be more effective for learning. Immediate feedback may encourage shallow processing (guess, check, adjust) rather than deep reflection (Butler et al., 2007). The efficiency of immediate correction may come at understanding cost.

Pedagogical Costs

There is a massive trade-off between optimising for efficiency. One is students might acquire superficial learning, that they just mastered the skill as a routine and then have poor adaptability in application (Chen et al., 2024). Also, there is a risk that continuous algorithmic support obstruct metacognitive skills such as ability for self-monitoring and planning to become underdeveloped among learners (Jansen et al., 2023). Further, removing obstacles too easily stifles resilience because students do not get chances to develop perseverance and self-efficacy with problem-solving. Also, efficient learning yields fast returns but also results in fragile knowledge that does not adapt well to new settings or withstands the test of time (Bjork & Yue, 2022).

Navigating the Tension

There are a number of strategic approaches to balancing efficiency and deeper learning. First, it is to add an element of deliberate difficulty by generating productive difficulties, such as solving problems before instruction (pretraining), spacing conditions and helping practitioners to focus on recognition over recall during practice (Kapur, 2022; Rohrer & Taylor, 2023). Second, support adaptive fading of scaffolds by gradually removing support as students become more proficient, so that they learn rather than depend on structure (Jansen et al., 2023). Third, make the efficiency-depth trade-off explicit to students so they realize that fast shortcuts sometimes come at a cost of understanding, and that challenging sequences help build lasting knowledge (Duan et al., 2024). Lastly, embrace optimism and a focus on the long term by impending systems that incentivise retention and transfer over time as opposed to current performance, moving from efficiency to more permanent understanding (Bjork & Bjork, 2022; Sense et al., 2023).

Tension 4: Individual Pathways vs. Collaborative Learning

Theoretical Tension

Computer-based personalisation at its roots makes learning personal through the diagnosis of knowledge states, by targeting specific needs, and by tracking individual growth. And yet this is at odds with sociocultural views that learning is by its nature social (Kalantzis & Cope, 2023; Schwartz et al., 2022). Knowledge is cultivated in dialogue, negotiation and participatory sense-making where learners produce ideas, are exposed to different perspectives, and co-develop understandings. As a matter of fact, social interaction is not only enriched as an add-on to this: it in essence structurally forms cognition itself. As a result, there is a tension: algorithms try to maximize the utility of individual pathways but they can damage group learning. As courses customise content, tempo, and level to individual learners more and more, students are exposed to different materials and representations thus reducing the common base that is needed for a meaningful collaboration.

Manifestations in Practice

Adaptive learning tools tailor personalized paths for individual learners so that each encounters appropriate challenge. However, this individualization brings up collaborative barriers: because students are working on completely different things whole class discussions cannot work (since not everyone has the same factors), and since many representations of a given idea are floating around peer-supported collaboration is nearly impossible. Furthermore, the awareness of common knowledge construction mechanisms is absent in these systems that consider group work to be potentially contaminating individual tests (Holstein et al. 2023).

AI tutoring chatbots like Khanmigo do the same with personalized explanations for each unique confusion. But that's privatized learning which could otherwise be social — as some other students are asking their question individually, they're getting individual answers without the opportunity to co-construct understanding in a group lesson. Thus, the chatbot becomes a one-person answer distributor instead of integrative collaboration (Kasneci et al., 2023).

Massive Open Online Courses (MOOCs) recommender systems revolve around the personalisation of courses that leverage users' backgrounds and goals while discovering content. However, this fragmentation is detrimental to community building. "Students are doing that less," he says, adding "guys' study sessions aren't as big, opportunistically." In the end, the platform consists of a series of disconnected learners rather than a community of practice (Reich & Ruipérez-Valiente, 2022).

Pedagogical Costs

Personalizing learning is very damaging to society. Firstly, the participants' level of atrophy is related to collaborative skills, since students have fewer opportunities to practice negotiating ideas and co-constructing knowledge both necessary in contemporary workplaces (Borge & Hod, 2022). Peer learning opportunities vanish as well; teaching others is one of the most powerful ways to understand, but personal systems do away with these opportunities (Fiorella & Mayer, 2023). In addition, competing perspectives are not transmitted when personalised pathways turn information into echo chambers, shielding one from competing views that can help learn better (Kapur & Bielaczyc, 2022). Lastly, community does not materialize as algorithmically-driven personalized systems elsewhere devoid of shared experiences required for learning communities that support student motivation, identity development and social connectedness) (Knight et al., 2023).

Navigating the Tension

Approaches for sustaining social learning enabling personalisation are:

Stratified individualisation: Individualise some of the dimensions of learning and standardise others in order to facilitate collaboration. Create personalized practice problems and pacing,

while standardizing conceptual introductions and project themes for example. In them is established common such that individual differences are respected (Stahl et al., 2014; Järvelä & Hadwin, 2013).

Collaborative personalisation: Don't personalize to the individual; personalize to small groups. Group up by knowledge/skill set, then cater content towards the group's collective need. This preserves benefits of personalisation while retaining social learning (Soller et al., 2005; Dillenbourg & Jermann, 2010).

Algorithmic collaboration: Deploy algorithms to enhance not substitute for collaborations. Systems may detect students with complementary knowledge for forming study groups, surface common confusions to foster discussion and recommend collaboration when multiple different students are experiencing difficulty on related topics. This renders algorithms as collaboration enablers, instead of replacements (Cukurova et al., 2018; Wise & Schwarz, 2017).

Compulsory collaboration spaces: Despite personalised pathways, insist on regular collaborative engagements which are not algorithmic-mediated. That there's time for student-led discussion, peer-teaching or group projects in which algorithms take a step back. This guarantees that social learning takes place even in algorithmically personalized systems (Topping, 2005; van der Pol et al., 2008).

Design Principles for Navigating Pedagogical Tensions

The preceding analysis reveals that algorithmic education involves navigating tensions rather than solving problems. Consequently, this section presents design principles for guiding this navigation.

Pedagogical Transparency

First, systems should make value hierarchies explicit. Algorithmic platforms embody choices about what matters in learning, and rather than obscuring these as neutral technical decisions, they should communicate them transparently. Systems that automate student recommendation should have clear priorities for when to promote efficiency over depth, or engagement over challenge, and leave educators and learners in an informed position to follow or disregard the system recommendations (Costanza-Chock et al.

In addition, systems should communicate trade-offs transparently. Users have to realize that enhancing one dimension (e.g., pathway efficiency) may penalize other aspects, shorter pathways are less deep (sacrificing depth for efficiency), or individualization of services can make individuals "more solitary" leaving less room for collaboration. This avoids the confusing of algorithmic recommendations with full-fledged pedagogical decisions (Horneber & Laumer, 2023).

Systems should also provide their reasoning, rather than deducing the black-box outputs. For instance: "This is your recommend problem as you have mastered the line and this will challenge you." These explanations assist users in assessing if the algorithmic reasoning is coherent with their objectives (Salminen et al., 2022).

Human-Algorithm Complementarity

The system needs to maintain the human ability to override. Teachers and students should still maintain the ability to defy recommendations, no matter how clever or not-so-clever you may think your algorithm is. Therefore, systems ought to serve as advisory tools that demand human judgment rather than prescriptive authorities, preserving educator professionalism and student agency (Holstein & McLaren, 2023). Also, design must delegate what algorithms can do to them, and keep in its domain what they cannot. Algorithms are great at digesting copious amounts of data and spotting patterns; humans are good at making sense of all types of contextual judgment,

nuance. In this way systems might take advantage of such complementary capabilities, allowing algorithms to perform pattern matching and information recall with the human role retained in interpreting struggles and rendering value judgments (Selwyn, 2023).

Lastly, algorithms should enable rather than replace collaboration. Rather than framing them as one-to-one tutors, design them to support group work and nudge students toward productive partner assignments or synthesizing class insights (Wise et al., 2022).

Bounded Algorithmic Scope

Systems should respect the algorithmic bounds by defining explicitly what they fail to measure: aesthetic appreciation, a personal sense of right and wrong, identity formation. Recognising these constraints helps avoid reifying algorithmic representations of learning (Farrow, 2023).

Furthermore, education should safeguard unoptimized time through defended opportunities for student-led investigation and open-ended projects explicitly beyond the pale of algorithmic evaluation. This guarantees that the aspects of experiences that algorithms cannot measure are put back to where they belong (Greene, 2021).

Finally, resist totalizing optimization. While algorithms might enhance some dimensions, others such as, wonder, curiosity, intellectual risk-taking, may be undermined by optimization. Therefore, maintain pedagogical commitment to values beyond algorithmic reach (Prinsloo & Slade, 2023).

Applying the PTF: A Case Analysis in Higher Education

To illustrate PTF in action, consider AI-powered tutoring in university settings. Khan Academy's Khanmigo is designed as a Socratic tutor: it challenges students to think critically and solve problems without giving direct answers. This design explicitly addresses Efficiency versus Deep Learning tension by slowing the pace to promote deeper reasoning. It aligns with pedagogical goals by guiding learners through hints, partially resolving Algorithmic Optimization versus Pedagogical Goals in favor of the latter. However, Khanmigo's individualized dialogue can heighten Individual Pathways versus Collaborative tension: students may become adept at one-on-one reasoning with AI but have fewer opportunities for peer discourse. Additionally, because content is organized around fixed mastery playlists, the system embodies tension between Standardization and Differentiation. The AI tutor differentiates within curriculum-aligned modules but cannot easily offer completely novel learning sequences outside standard content.

By contrast, models like ChatGPT typically prioritize efficiency: delivering direct answers almost instantly. This maximizes efficiency and student satisfaction short-term but tends to sacrifice deep cognitive processing. ChatGPT's default mode intensifies Efficiency versus Deep Learning tension (students get quick solutions rather than grapple with problems). Recent studies in higher education (Neumann et al., 2025; Chan et al., 2023) note that unlike Khanmigo, ChatGPT doesn't guide learners to find answers themselves. This difference shows that ChatGPT-like tools optimize for immediate performance, potentially at the expense of meaningful learning. Furthermore, using ChatGPT is highly individual, exacerbating Individual versus Collaborative tension by reducing incentives to work together.

These examples demonstrate how applying PTF reveals specific trade-offs in higher education contexts: Khanmigo's design sacrifices some efficiency to bolster conceptual understanding, whereas ChatGPT maximizes efficiency but risks superficial learning. In each case, PTF helps identify which tensions are emphasized and what pedagogical values are compromised. Educators and designers can make these trade-offs explicit (for instance, by integrating collaborative activities alongside personalized AI tutoring) to mitigate unintended consequences in university settings.

Implications and Conclusion

Implications for Research

This will require future research to explore how such tensions play out in different educational contexts, cultures and disciplines while creating measures to evaluate pedagogical trade-offs. In addition to replication studies, future work could investigate long-term effects on metacognitive and collaborating skills and intrinsic motivation.

Implications for Practice

For instructors, the framework offers modes of analysis for assessing technologies by questioning: “What is this optimising for? What values does it prioritise? What dimensions does it neglect?”

For designers, the framework provides support in the design of ethically-informed development for designers. Instead of taking for granted that product advice always benefits from personalisation, determine which dimensions are assisted by algorithmic mediation and work to design transparency mechanisms that make it easier for users to recognise these limitations.

For policymakers, the framework indicates considerations for evaluation beyond efficiency to consider how systems form learning experiences and preserve educator professionals.

Limitations

This framework emphasizes conflict, and probably underemphasizes the places where real agreement can be found. It also focuses on the content created algorithms, not scoring systems and it comes from Western education theories. 8. 4 Conclusion

The promise of algorithmic thinking is still highly attractive, but realizing it necessitates coming to terms with trade-offs that need to be made. Instead of searching for perfect systems, we should aim for systems that reveal their value hierarchies and leaves space for barely precise outcomes with algorithms as pedagogical devices that need to be thoughtfully implemented.

The role of education isn't merely to optimize toward an algorithm. Algorithms may be efficient, but they are not a substitute for human judgement about what counts as significant learning. The question, then, isn't whether to use algorithmic systems but how to use them well, in a way that brings out tensions into the open, exposes trade-offs and makes pedagogical values central.

Scope and Applicability in Higher Education

PTF is intended for analyzing content-oriented educational algorithms—systems that select, sequence, or generate learning materials (e.g., intelligent tutors, adaptive quizzes, generative AI tutors). We deliberately exclude algorithmic systems whose primary function is assessment, student risk prediction, or administrative decision-making. The framework is broadly applicable across higher education, drawing on foundational learning theories (constructivist, socio-cultural, critical) that span educational stages.

Higher education contexts differ from K-12/secondary education: university settings with discipline-specific requirements may experience Standardization versus Differentiation tension differently than K-12/secondary education with mandated standards. In higher education, Efficiency versus Deep Learning tension manifests differently in massive online courses balancing completion rates against conceptual mastery. Recent research on personalized adaptive learning in higher education (Naseer et al., 2024; Ma et al., 2023) demonstrates these tensions emerge distinctly at university level, where students have greater autonomy but face complex disciplinary knowledge.

Subject domain affects how tensions play out. In well-structured domains like mathematics or language exercises, algorithmic objectives are easier to define, potentially intensifying Algorithmic Optimization versus Pedagogical Goals. In creative or open-ended domains (e.g.,

creative writing, art), pedagogical goals like originality are harder to measure, so algorithmic systems tend to impose stronger constraints, heightening Standardization versus Differentiation tension. Studies on AI-driven adaptive learning (Adiguzel et al., 2024; Demartini et al., 2024) show these domain-specific challenges are particularly acute in higher education's diverse disciplinary contexts.

Algorithm type influences applicability. Traditional rule-based systems allow designers to map tensions explicitly, whereas modern ML-based or large language models offer greater adaptability but less interpretability. Neural models can exacerbate the "black-box" aspect of Algorithmic Optimization—making it harder for educators to align algorithm's internal logic with pedagogical intent. Recent higher education implementations (Fischer & Dobbins, 2024) demonstrate that neural adaptive systems might relieve faculty of design work but create new tensions around transparency and trust. PTF's tensions remain relevant regardless of algorithmic method; the difference is in implementation.

In summary, PTF is a flexible analytic lens valid across higher education disciplines and contexts, but specific implications vary. We encourage researchers and practitioners to consider contextual factors—discipline, subject, and algorithmic complexity—when using the framework, recognizing that some pedagogical values may carry different weights in different university settings.

References

- Adiguzel, T., de Vries, B., & Jing, L. (2024). AI-driven adaptive learning for sustainable educational transformation. *Sustainable Development*, 32(5), 4296–4310. <https://doi.org/10.1002/sd.3221>
- Baker, R. S. (2022). Algorithmic bias in education. *International Journal of Artificial Intelligence in Education*, 32(4), 1052–1092. <https://doi.org/10.1007/s40593-021-00285-9>
- Bjork, R. A., & Yue, C. L. (2022). Desirable difficulties in learning: Theory and applications. In K. R. Sawyer (Ed.), *The Cambridge handbook of the learning sciences* (3rd ed., pp. 113–126). Cambridge University Press.
- Bohunovsky, L., Radinger-Peer, V., Zint, M., & Penker, M. (2023). Change agents under tensions: A paradox approach to strategies for transforming higher education toward sustainability. *International Journal of Sustainability in Higher Education*, 24(9), 372–392. <https://doi.org/10.1108/IJSHE-12-2022-0393>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(1), Article 43. <https://doi.org/10.1186/s41239-023-00411-8>
- Chen, L., Zhang, Y., & Wang, X. (2024). Procedural fluency versus conceptual understanding: The hidden costs of algorithmic instruction in mathematics education. *Journal of Educational Psychology*, 116(2), 245–261.
- Chowrira, S. G., Smith, K. M., Dubois, M., Roche Allred, Z., & Wieman, C. E. (2023). DIY productive failure: Boosting performance in a large undergraduate biology course. *npj Science of Learning*, 8, Article 7. <https://doi.org/10.1038/s41539-023-00156-2>
- Demartini, C. G., Sciascia, L., Bosso, A., & Manuri, F. (2024). Artificial intelligence bringing improvements to adaptive learning in education: A case study. *Sustainability*, 16(4), Article 1347. <https://doi.org/10.3390/su16041347>

- Dziuban, C., Moskal, P., Parker, L., Campbell, M., Howlin, C., & Johnson, C. (2016). Adaptive learning: A stabilizing influence across disciplines and universities. *Online Learning Journal*, 20(3). <https://doi.org/10.24059/olj.v20i3.1074>
- Favero, L., Pérez-Ortiz, M., & Manrique, D. (2024). Enhancing critical thinking in education by means of a Socratic chatbot. *Expert Systems with Applications*, 246, Article 123132. <https://doi.org/10.1016/j.eswa.2024.123132>
- Fischer, I., & Dobbins, K. (2024). Is it worth it? How paradoxical tensions of identity shape the readiness of management educators to embrace transformative technologies in their teaching. *Journal of Management Education*, 48(2), 234–264. <https://doi.org/10.1177/10525629231201843>
- Freire, P. (1970). *Pedagogy of the oppressed* (M. B. Ramos, Trans.). Continuum.
- Gallagher, M., & Breines, M. (2022). Unpacking the hidden curricula in educational automation: A methodology for ethical praxis. *Postdigital Science and Education*. <https://doi.org/10.1007/s42438-022-00342-z>
- Hall, R., Bowden, M., & Thorogood, J. (2025). Teaching in the AI era: Sustainable digital education through ethical integration and teacher empowerment. *Sustainability*, 17(16), Article 7405. <https://doi.org/10.3390/su17167405>
- Hazzan-Bishara, A., Kol, O., & Levy, S. (2025). The factors affecting teachers' adoption of AI technologies: A unified model of external and internal determinants. *Education and Information Technologies*, 30, 15043–15069. <https://doi.org/10.1007/s10639-025-13393-z>
- Holstein, K., Aleven, V., & Rummel, N. (2023). Co-designing adaptive learning technologies with teachers: Balancing personalisation and classroom orchestration. *British Journal of Educational Technology*, 54(1), 223–245.
- Hwang, S., Suzuki, Y., & Saito, K. (2024). Undesirable difficulty of interleaved practice: The importance of initial blocked practice for declarative knowledge development in low-achieving adolescents. *Language Learning*, 74(3), 689–736. <https://doi.org/10.1111/lang.12659>
- ILO Group. (2024). Framework for implementing artificial intelligence (AI) in K-12 education v1.0. https://www.ilogroup.com/wp-content/uploads/2024/03/Framework-for-Implementing-Artificial-Intelligence-AI-in-K-12-Education_v1.0.pdf
- Jansen, R. S., van Leeuwen, A., & Janssen, J. (2023). Metacognitive scaffolding and self-regulated learning in digital environments: When support becomes dependency. *Computers & Education*, 195, Article 104725.
- Kasneci, E., Seßler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, Article 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kemp, N., & Scoffham, S. (2022). The paradox model: Towards a conceptual framework for engaging with sustainability in higher education. *Sustainability*, 14(19), Article 12285. <https://doi.org/10.3390/su141912285>
- Khan Academy. (n.d.). Meet Khanmigo: Khan Academy's AI-powered teaching assistant & tutor. Retrieved November 2025, from <https://www.khanmigo.ai/>

- Lane, S. (2020). Test-based accountability systems: The importance of paying attention to consequences. ETS Research Report Series, 2020(1), 1–16. <https://doi.org/10.1002/ets2.12283>
- Ma, Y., Wang, L., Zhang, J., Liu, F., & Jiang, Q. (2023). A personalized learning path recommendation method incorporating multi-algorithm. *Applied Sciences*, 13(10), Article 5946. <https://doi.org/10.3390/app13105946>
- Naseer, F., Khan, M. N., Tahir, M., Addas, A., & Aejaz, S. M. H. (2024). Integrating deep learning techniques for personalized learning pathways in higher education. *Heliyon*, 10(12), Article e32628. <https://doi.org/10.1016/j.heliyon.2024.e32628>
- Nelson, A., & Elias, K. L. (2023). Desirable difficulty: Theory and application of intentionally challenging learning. *Medical Education*, 57(2), 123–130. <https://doi.org/10.1111/medu.14916>
- Neumann, A. T., Yin, Y., Sowe, S., Decker, S., & Jarke, M. (2025). An LLM-driven chatbot in higher education for databases and information systems. *IEEE Transactions on Education*, 68(1), 103–116. <https://doi.org/10.1109/TE.2024.3425707>
- Pelánek, R. (2024). Adaptive learning is hard: Challenges, nuances, and trade-offs in modeling. *International Journal of Artificial Intelligence in Education*, 34, 892–915. <https://doi.org/10.1007/s40593-024-00400-6>
- Peng, H., Ma, S., & Spector, J. M. (2019). Personalized adaptive learning: An emerging pedagogical approach enabled by a smart learning environment. *Smart Learning Environments*, 6(1), Article 9. <https://doi.org/10.1186/s40561-019-0089-y>
- Piaget, J. (1952). *The origins of intelligence in children* (M. Cook, Trans.). International Universities Press.
- Richter, F. (2025). *The ethical acceptability of personalisation via intelligent systems in education* [Doctoral dissertation, Katholische Universität Eichstätt-Ingolstadt]. <https://doi.org/10.1007/s00146-025-02690-3>
- Sangwa, N., Nyandoro, F., & Dlamini, T. (2025). Dynamic AI governance in education (DAIGE): A multi-level framework. *Educational Technology Research and Development*. Advance online publication.
- Schaap, L., van der Heijden, B., & Korzilius, H. (2024). Why so many change efforts fail: Using paradox theory as a lens to understand the complexity of educational change. *International Journal of Leadership in Education*. Advance online publication. <https://doi.org/10.1080/13603124.2023.2298210>
- Skinner, B. F. (1958). Teaching machines. *Science*, 128(3330), 969–977. <https://doi.org/10.1126/science.128.3330.969>
- Steenhof, N., Woods, N. N., & Mylopoulos, M. (2023). Adaptive expertise in undergraduate pharmacy education. *Pharmacy*, 11(1), Article 32. <https://doi.org/10.3390/pharmacy11010032>
- Strielkowski, W., Zenchenko, S., Tarasova, A., & Radyukova, Y. (2024). AI-driven adaptive learning for sustainable educational transformation. *Sustainable Development*, 32(5), 4296–4310. <https://doi.org/10.1002/sd.3221>
- Thrush, T., Jiang, J., Bartolo, M., Nie, A., Kaushik, D., Jiang, C., Tang, T., Kann, K., Potts, C., & Williams, A. (2024). The language model bias database. <https://arxiv.org/abs/2403.18915>

- UNESCO. (2024). AI competency framework for teachers. https://www.cedefop.europa.eu/files/unesco_ai_competency_framework_for_teachers.pdf
- Verghis, A. M., Jose, B., Sheeja, P. R., Varghese, S. M., Mumthas, S., & Shaiju, K. S. (2025). Beyond personalisation: Autonomy and agency in intelligent systems education. *Frontiers in Education*, 10, Article 1610239. <https://doi.org/10.3389/feduc.2025.1610239>
- Vincent-Lancrin, S., González-Sancho, C., Bouckaert, M., de Luca, F., Fernández-Barrerra, M., Jacotin, G., Urgel, J., & Vidal, Q. (2019). Fostering students' creativity and critical thinking: What it means in school. OECD Publishing. <https://doi.org/10.1787/62212c37-en>
- Xu, X., Li, J., Zhao, Y., & Chen, Q. (2025). AI optimization algorithms enhance higher education management and personalized teaching through empirical analysis. *Scientific Reports*, 15, Article 94481. <https://doi.org/10.1038/s41598-025-94481-5>
- Yang, F., Huang, R., Wu, C., Liu, D., & Zheng, L. (2025). A framework for AI ethics literacy: Development, validation, and its role in fostering students' self-rated learning competence. *Scientific Reports*, 15, Article 21977. <https://doi.org/10.1038/s41598-025-21977-5>
- Zhou, T., Sun, X., & Wang, P. (2025). Algorithmic optimisation in AI-driven education: Balancing metrics and meaning. *Computers & Education: AI*, 6, Article 100212.