Natural Resources for Human Health
eISSN: 2583-1194

Original Research

**Received** 16 April 2026
**Revised** 18 May 2026
**Accepted** 20 June 2026

Natural Resources for Human Health 2026, Vol. 6 (5s): 1030–1038

**KEYWORDS**:
Digital Transformation, Artificial Intelligence, Education Technology, Legal Technology, Learning Analytics, Algorithmic Accountability, AI Governance, Legal Compliance.

# Digital Transformation in Education & Law: Leveraging Artificial Intelligence for Learning Innovation, Legal Compliance, and Governance Excellence

Dr. Arpita Sehgal[1], Dr. S M Abdul Mannan Hussain[2], Budati Subbaiah[3], Dr. Pratap Kumar Sahu[4], Lonavath Srinivas Naik[5], Dr. Mitali Gupta[6]

[1]Assistant Professor, CLS, GIBs, North Delhi, arpitasadvocate@gmail.com
[2]Associate Professor, Project Management, Regenesys Business School, Sandton, Johannesburg, Gauteng, South Africa, sohailkhn248@gmail.com
[3]Assistant professor, Department of Computer science, Kakaraparti Bhavanarayana College Autonomous Vijayawada, NTR District, Andhra Pradesh, subbu.epm@gmail.com
[4]Assistant Professor, Department of Bachelor of Business Administration, Parul Institute of Business Administration, Faculty of Management Studies, Vadodara, Gujarat.pks26061983@gmail.com, ORCID:0009-0001-4198-7898
[5]Assistant Professor, Department of Mechanical Engineering, School of Engineering, Anurag University, Medchal Malkajgiri, Hyderabad, Telangana , srinivasnaikmech@anurag.edu.in
[6]Assistant professor, Daimsr, Nagpur, mitali_gupta@daimsr.edu.in

**ABSTRACT:** Education and law are institutionally distinct sectors, yet both are undergoing a structurally similar digital transformation: artificial intelligence (AI) is being adopted simultaneously as an innovation tool that improves core professional practice, personalized instruction in education, predictive and analytical support in law, and as a governance object requiring new compliance, privacy, and accountability frameworks precisely because of its adoption. This paper reviews the AI-driven digital transformation literature across both sectors within a single comparative framework, synthesizing intelligent tutoring systems and learning analytics research in education with legal analytics, predictive justice, and legal-technology research in law, before examining the governance, data-privacy, and algorithmic-accountability literature that has emerged in response to AI adoption in both domains. Particular attention is given to the structural parallel between the two sectors: both involve consequential, sometimes life-altering decisions about individuals, students and litigants respectively, made partly through AI-derived predictions, and both have accordingly developed parallel, though not identical, due-process and explicability concerns. Comparative tables map AI application types onto their core function across education and law, cross-reference governance and compliance concerns against the specific mechanism each addresses, and set application maturity against evidentiary support in both sectors. The paper concludes that digital transformation in both education and law has advanced learning and legal-practice innovation considerably faster than the governance infrastructure needed to ensure these AI applications operate fairly, transparently, and in compliance with existing legal and regulatory obligations, and identifies cross-sector governance harmonization as the central future research prospect.

## 1. INTRODUCTION

Artificial intelligence has been adopted across education and law along a structurally parallel trajectory: both sectors have embraced AI-driven tools intended to improve the core professional or instructional task, personalized teaching and legal analysis respectively, while simultaneously confronting governance questions the underlying professions' existing regulatory frameworks were not originally designed to address [1], [6]. Zawacki-Richter, Marín, Bond, and Gouverneur's systematic review of AI applications in higher education found that the field had grown rapidly across four principal application clusters, profiling and prediction, intelligent tutoring systems, assessment and evaluation, and adaptive systems and personalization, while also finding that the reviewed literature disproportionately emphasized technical implementation over pedagogical, ethical, or educational-theory grounding, a gap this paper's governance-focused discussion addresses directly [1]. VanLehn's meta-analytic comparison of intelligent tutoring systems against human tutoring and no-tutoring control conditions found that well-designed intelligent tutoring systems achieved learning gains approaching, though not fully matching, those produced by expert human tutors, and substantially exceeding no-tutoring control conditions, establishing a robust empirical foundation for AI-driven personalized instruction specifically rather than treating it as a purely aspirational technology [2].

Siemens's foundational framing of learning analytics as an emerging academic discipline argued that the large-scale behavioral and performance data generated by digital learning environments could be systematically analyzed to predict student outcomes, identify at-risk learners, and inform pedagogical intervention, a framing that positioned learning analytics as a distinct but closely related strand of educational AI alongside intelligent tutoring systems specifically [3]. Holmes, Bialik, and Fadel's comprehensive treatment of AI in education extended this analysis to the full range of application types, arguing that AI's educational promise depends critically on how deliberately these systems are designed around established learning-science principles rather than around technical capability alone, a caution that anticipates this paper's central concern with the gap between AI capability and governance readiness [4]. Selwyn's more critical treatment of AI in education cautioned specifically against uncritical technological optimism, arguing that AI-driven personalization and prediction systems in education carry embedded assumptions about what constitutes good learning and a good learner that are rarely made explicit or contested, a concern with direct relevance to the algorithmic-accountability literature this paper reviews later for both education and law [7].

The legal profession has undergone a parallel, though independently developed, AI-adoption trajectory. Susskind's influential analysis of the legal profession's technological future argued that AI and related technologies would progressively decompose legal work into discrete tasks, document review, legal research, contract analysis, and dispute prediction among them, some of which are highly amenable to automation even where the practice of law as a whole is not, a task-decomposition argument structurally similar to the prediction-machine framing subsequently developed for AI-driven business decision-making more broadly [8]. Surden's foundational treatment of machine learning and law examined this task-decomposition argument from a more technical legal-scholarship perspective, arguing that machine learning's comparative advantage in law lies specifically in pattern recognition across large volumes of legal text and case outcomes, a capability well suited to predictive analytics and legal research automation but poorly suited to the normative reasoning and judgment genuine legal argumentation requires [9]. Katz, Bommarito, and Blackman's widely cited empirical demonstration of machine learning-based prediction of U.S. Supreme Court case outcomes, achieving prediction accuracy exceeding chance and comparable to or exceeding some expert-based prediction approaches, provided direct empirical evidence for legal AI's predictive capability, while simultaneously raising the governance question of what it means for judicial outcomes to be statistically predictable from case features alone [10].

This predictive capability, applied not merely to academic case-outcome prediction but to actual judicial and administrative decision-making, has generated the legal sector's most consequential governance controversy. Angwin, Larson, Mattu, and Kirchner's investigative analysis of the COMPAS recidivism-risk-assessment algorithm, used by U.S. courts to inform bail and sentencing decisions, found that the algorithm's risk classifications exhibited a racially disparate error pattern, with Black defendants more likely to be incorrectly classified as high-risk relative to white defendants, a finding that catalyzed a substantial subsequent academic and policy debate over algorithmic accountability in judicial decision-making specifically [13]. Citron's earlier, foundational legal analysis of "technological due process" argued that automated decision-making systems used in government and judicial contexts require dedicated due-process protections, including transparency about the system's decision logic and a meaningful opportunity to contest an automated determination, that existing administrative and constitutional law doctrine, developed for human decision-makers, does not automatically provide, a framework directly applicable to both the judicial AI concerns Angwin et al. document and the algorithmic decision-making increasingly present in educational admissions, grading, and disciplinary contexts [11].

This paper reviews the intelligent tutoring, learning analytics, legal analytics, and predictive-justice literatures within a single comparative framework, before examining the governance, data-privacy, and algorithmic-accountability literature that has developed, largely in parallel, across both sectors in response to this AI adoption.

**Research Objectives**

- To review the AI-driven learning innovation literature, including intelligent tutoring systems and learning analytics, in education.
- To synthesize the legal analytics and predictive-justice literature documenting AI adoption in legal practice and judicial decision-making.
- To examine the governance, data-privacy, and algorithmic-accountability literature developed in response to AI adoption in both sectors.
- To evaluate the structural parallels and divergences between education and law's AI-governance challenges.
- To identify future research priorities for harmonizing AI governance frameworks across education and legal institutions.

## 2. LITERATURE REVIEW

### 2.1 Intelligent Tutoring Systems and Personalized Learning

Woolf's comprehensive treatment of intelligent interactive tutoring systems formalized the field's standard architecture, comprising a domain model representing the subject-matter expertise being taught, a student model representing an individual learner's current knowledge state, and a pedagogical model determining what instructional action to take given the gap between the two, an architecture that has remained the conceptual backbone of subsequent intelligent tutoring system development [5]. VanLehn's meta-analysis, synthesizing effect sizes across a large number of controlled studies comparing intelligent tutoring systems against human tutoring, classroom instruction, and no-treatment control conditions, found that step-based intelligent tutoring systems, providing feedback on each individual problem-solving step, achieved learning gains statistically indistinguishable from human tutoring in several of the reviewed studies, while answer-based systems, providing feedback only on final answers, achieved smaller though still educationally meaningful gains, establishing a specific design distinction with direct implications for how intelligent tutoring systems should be built to maximize learning benefit [2].

### 2.2 Learning Analytics and Predictive Educational Systems

Siemens's foundational learning-analytics framing distinguished the discipline from the closely related field of educational data mining primarily by emphasis, with learning analytics oriented more toward informing human decision-making, by instructors, administrators, and learners themselves, and educational data mining oriented more toward automated discovery of patterns and automated system adaptation, a distinction that remains useful for situating where human judgment versus automated action is expected to operate within a given educational AI system [3]. Baker and Inventado's complementary treatment of educational data mining and learning analytics catalogued specific predictive application types, including early-warning systems identifying students at risk of course failure or dropout, and found consistent evidence across the reviewed studies that such systems could identify at-risk students with useful lead time for intervention, while cautioning that predictive accuracy alone does not guarantee that the resulting human intervention is effective, timely, or appropriately targeted [6].

### 2.3 Legal Analytics and Predictive Justice

Surden's treatment of machine learning and law characterized legal AI's core technical capability as pattern extraction from large corpora of legal text, case law, statutes, and contracts, arguing that this capability supports legal research automation, contract review, and outcome prediction considerably more directly than it supports the normative, values-laden reasoning legal argumentation and judicial decision-making require [9]. Katz, Bommarito, and Blackman's empirical Supreme Court prediction study applied a random-forest machine learning model to a large historical dataset of case features and outcomes, achieving case-outcome prediction accuracy and individual-justice-vote prediction accuracy both exceeding relevant baseline comparisons, and the authors specifically situated this result as evidence that substantial, previously implicit structure in judicial decision-making can be extracted through systematic machine learning analysis of historical case data [10]. Remus and Levy's more skeptical, practice-oriented analysis of whether "robots can be lawyers" examined AI adoption across specific legal-practice task categories, document review, legal research, and contract drafting particularly, finding that AI adoption had already substantially automated portions of these specific tasks, while arguing that the broader practice of law, requiring client counseling, courtroom advocacy, and complex legal judgment, remained considerably more resistant to automation than either optimistic or alarmist popular accounts of "robot lawyers" suggested [12]. Ashley's comprehensive treatment of artificial intelligence and legal analytics extended this task-specific analysis into a broader argument that legal AI's most durable contribution lies in argument-mining and case-based reasoning support tools that assist, rather than replace, human legal reasoning, a framing structurally parallel to Woolf's and VanLehn's intelligent-tutoring framing of AI as augmenting rather than replacing human instruction [2], [5].

**2.4 Algorithmic Accountability in Judicial and Administrative Decision-Making**

Angwin, Larson, Mattu, and Kirchner's ProPublica investigation of the COMPAS recidivism-risk algorithm found a specific, racially patterned error asymmetry, Black defendants who did not subsequently reoffend were more likely to be misclassified as high-risk than white defendants with the same actual outcome, and white defendants who did subsequently reoffend were more likely to be misclassified as low-risk than Black defendants with the same actual outcome, a finding that catalyzed extensive subsequent academic debate over how to define and measure algorithmic fairness in exactly this kind of consequential judicial decision-support context [13]. Citron's technological due process framework, developed prior to this specific controversy but directly applicable to it, argued that automated government decision-making systems require transparency, contestability, and accuracy-auditing safeguards analogous to the due-process protections traditionally afforded to human administrative decision-making, safeguards Citron argued were routinely absent from early automated benefits-determination and, by extension, risk-assessment systems [11]. Re and Solow-Niederman's subsequent legal-scholarship analysis of "artificially intelligent justice" extended this due-process concern specifically to judicial AI applications, distinguishing AI tools that merely support human judicial decision-making from those that substantively influence or replace it, and arguing that the appropriate governance response depends substantially on which of these two roles a given legal AI application actually occupies in practice, a distinction directly analogous to the human-control-versus-automation typology developed in the organizational AI governance literature more broadly [14].

**2.5 Governance, Transparency, and Ethical Frameworks Applicable Across Both Sectors**

Floridi et al.'s AI4People framework articulated beneficence, non-maleficence, autonomy, justice, and explicability as cross-domain ethical principles for responsible AI deployment, with explicability, the capacity to understand and articulate why an AI system produced a given output, carrying particular relevance to both the judicial-decision context Angwin et al. and Re and Solow-Niederman examine and the educational-assessment context where an AI-generated grade or admissions recommendation likewise demands justification [15]. Pasquale's broader analysis of algorithmic opacity across governmental, financial, and other socially consequential systems characterized this opacity as a structural feature of contemporary automated decision-making generally, arguing that the proprietary and technical complexity of many deployed algorithms creates a substantial accountability gap that neither existing sector-specific regulation, financial, educational, or judicial, was designed to close, a cross-sector framing directly relevant to why this paper treats education and law's governance challenges as structurally parallel rather than sector-specific curiosities [16]. Selbst and Barocas's legal analysis of explainable-machine-learning approaches specifically examined whether post-hoc explanation techniques satisfy the kind of legal explicability standard due-process and anti-discrimination law might require, finding that many technical explanation methods provide a plausible-sounding justification without necessarily providing an accurate account of the model's actual decision process, a concern directly paralleling Rudin's broader critique of post-hoc explanation for high-stakes decisions and carrying equal force whether the high-stakes decision in question is a parole determination or a high-stakes educational placement decision [17].

### 2.6 Data Privacy Governance Specific to Education

Zeide's legal analysis of the structural consequences of big-data-driven education argued that the large-scale student data collection underlying contemporary learning analytics and adaptive learning systems creates governance challenges that existing student-privacy law, developed originally for paper-based educational records, does not straightforwardly address, particularly regarding data collected by third-party educational-technology vendors operating outside traditional school governance structures [18]. Regan and Jesse's complementary ethical analysis of educational-technology data practices catalogued specific privacy risks, including the potential for granular behavioral data collected for legitimate pedagogical purposes to be repurposed for commercial or unrelated institutional uses, and argued that meaningful student and parental consent is difficult to obtain in practice given the complexity of contemporary educational-technology data flows, a practical consent challenge structurally similar to the informed-consent challenges long documented in other data-intensive sectors [19].

### 2.7 Research Gap

While the intelligent tutoring and learning analytics literature reviewed in Sections 2.1 and 2.2 has established robust evidence for AI's educational effectiveness, and the legal analytics and predictive-justice literature reviewed in Section 2.3 has established comparably robust evidence for AI's legal-practice utility, comparatively few studies directly compare the governance and accountability frameworks each sector has developed in response, despite the structural parallel Citron's and Re and Solow-Niederman's due-process analyses and Zeide's and Regan and Jesse's privacy analyses collectively suggest [11], [14], [18], [19]. This review addresses this gap by organizing the reviewed evidence within a comparative framework spanning application type, sector, and governance mechanism.

## 3. METHODOLOGY

This study adopts a qualitative, descriptive, and comparative research design synthesizing peer-reviewed education-technology, legal-technology, and AI-governance literature. Reviewed works were compared along four axes: sector (education or law); application type (instructional/analytical support versus decision-support/prediction); evidence type (controlled effectiveness study, empirical prediction study, or governance/legal analysis); and whether the work addresses innovation, compliance/governance, or both jointly.

## IV. Results and Discussion

### 4.1 AI Application Types Mapped to Core Function Across Education and Law

Table 1 maps specific AI application types onto their core function within each sector, a structure that makes the cross-sector structural parallel explicit rather than treating education and law as unrelated case studies.

**Table 1.** AI Application Types and Core Function in Education and Law

| Application Type | Sector | Core Function | Documented Effectiveness | Key Reference |
|---|---|---|---|---|
| Intelligent tutoring systems | Education | Adaptive, step-based instructional feedback | Learning gains approaching human-tutor levels (step-based systems) | VanLehn [2]; Woolf [5] |
| Learning analytics / early-warning systems | Education | Predicting at-risk students for timely intervention | Useful predictive lead time; intervention effectiveness less established | Siemens [3]; Baker & Inventado [6] |
| Legal research and document-review automation | Law | Pattern extraction from large legal text corpora | Substantial task-level automation achieved | Surden [9]; Remus & Levy [12] |
| Case-outcome / judicial-vote prediction | Law | Predicting court decisions from case features | Prediction accuracy exceeding relevant baselines | Katz, Bommarito & Blackman [10] |

| | | | | |
|---|---|---|---|---|
| Recidivism/risk-assessment algorithms | Law | Informing bail/sentencing decisions | Effective at aggregate prediction; documented disparate error rates | Angwin, Larson, Mattu & Kirchner [13] |
| Argument-mining / case-based reasoning support | Law | Assisting, not replacing, legal reasoning | Supports rather than substitutes for human judgment | Ashley [2] |

### 4.2 Governance and Compliance Concerns Cross-Referenced Against Mechanism

Table 2 shifts from application type to governance concern, cross-referencing each concern against the sector in which it has been most directly documented and the specific mechanism proposed to address it.

**Table 2.** Governance and Compliance Concerns Cross-Referenced Against Mechanism

| Governance Concern | Sector(s) Documented | Underlying Mechanism | Proposed/Documented Governance Response | Key Reference |
|---|---|---|---|---|
| Algorithmic due process | Law (primary); Education (analogous) | Automated decisions lack contestability/transparency safeguards | Technological due process framework (transparency, contestability, accuracy auditing) | Citron [11] |
| Disparate error rates across demographic groups | Law | Risk algorithms miscalibrated differently across groups | Fairness-metric auditing; algorithmic impact assessment | Angwin, Larson, Mattu & Kirchner [13] |
| Post-hoc explanation inadequacy | Law and Education (cross-cutting) | Explanations approximate rather than describe true model reasoning | Preference for inherently interpretable models in high-stakes contexts | Selbst & Barocas [17] |
| Algorithmic opacity generally | Cross-sector | Proprietary/technical complexity obscures decision logic | Explicability as an ethical design principle | Floridi et al. [15]; Pasquale [16] |
| Human-role ambiguity (support vs. substitution) | Law | Unclear whether AI supports or substitutes judicial judgment | Explicit classification of AI's decision-role (support vs. substantive influence) | Re & Solow-Niederman [14] |
| Student data privacy and third-party vendor governance | Education | Ed-tech vendors collect data outside traditional school governance | Updated student-privacy legal frameworks; vendor accountability requirements | Zeide [18] |
| Meaningful consent for data collection | Education | Complexity of data flows undermines informed consent | Simplified, genuinely informed consent mechanisms | Regan & Jesse [19] |

### 4.3 Application Maturity Against Evidentiary Support

Table 3 departs from Tables 1 and 2's technology- and governance-centered structures to assess each major application category's maturity, how far it has progressed from research demonstration toward routine deployment, against the strength of its supporting evidence base.

**Table 3.** Application Maturity Against Evidentiary Support

| Application | Sector | Maturity Stage | Evidentiary Support | Governance Readiness |
|---|---|---|---|---|
| Step-based intelligent tutoring systems | Education | Widely deployed in specific subject domains (math, science) | Strong (meta-analytic effect sizes) | Moderate (pedagogical governance more developed than data governance) |
| Early-warning/dropout-prediction systems | Education | Deployed at many institutions | Moderate (predictive accuracy established; intervention efficacy less so) | Weak (privacy/vendor governance lagging) |
| Legal research/document-review AI | Law | Widely deployed in practice | Strong (well-documented task automation) | Moderate (professional-responsibility rules adapting) |
| Judicial risk-assessment algorithms | Law | Deployed in multiple U.S. jurisdictions | Strong predictive accuracy; contested fairness properties | Weak (due-process and fairness standards unresolved) |
| Case-outcome prediction (research/academic) | Law | Largely research-stage | Strong within specific study designs | Not yet directly applicable (not deployed operationally) |
| Cross-sector algorithmic accountability standards | Both | Emerging; largely proposed rather than implemented | Conceptual frameworks well developed; empirical validation limited | Weak (no harmonized standard across sectors) |

## DISCUSSION

Read together, Tables 1 through 3 support this paper's central claim that AI-driven innovation in both education and law has substantially outpaced the governance infrastructure needed to ensure these applications operate fairly and transparently. Table 1 establishes that both sectors have achieved genuine, evidence-supported innovation, VanLehn's meta-analytic tutoring effect sizes and Katz, Bommarito, and Blackman's judicial-prediction accuracy each represent methodologically rigorous demonstrations of AI capability within their respective domains [2], [10]. Table 2's governance mapping, however, indicates that the specific accountability mechanisms each sector has developed in response remain considerably less mature and, notably, structurally similar across sectors despite having developed largely independently: Citron's due-process framework and Zeide's student-privacy framework both identify a governance gap between existing regulatory frameworks, designed for human decision-makers or paper-based records, and the automated, data-intensive systems now operating within their jurisdiction [11], [18]. Table 3's maturity-versus-readiness assessment makes this gap concrete, showing that the applications with the strongest deployment maturity, judicial risk-assessment algorithms and legal research automation, are not the applications with the strongest governance readiness, a mismatch Angwin et al.'s COMPAS findings illustrate directly for the judicial context and that Regan and Jesse's consent-adequacy concerns illustrate for the educational context [11], [13], [18], [19]. This pattern suggests that both sectors face the same underlying structural problem Pasquale identifies at a cross-domain level, algorithmic opacity outpacing existing accountability mechanisms, rather than two unrelated, sector-specific governance challenges, a reading that directly supports this paper's argument for treating education and law's AI governance needs as a shared research problem rather than two separately theorized literatures [15], [16], [17].

## CONCLUSION

This paper has reviewed the AI-driven digital transformation literature across education and law within a single comparative framework, synthesizing intelligent tutoring systems and learning analytics research with legal analytics and predictive-justice research, before examining the governance, privacy, and algorithmic-accountability literature that has developed in response to AI adoption in both sectors. The evidence indicates that both sectors have achieved genuine, well-evidenced innovation, AI-driven personalized instruction achieving learning gains approaching expert human tutoring, and legal AI achieving substantial task automation and meaningful predictive accuracy for judicial outcomes, while simultaneously confronting a structurally similar governance gap between existing regulatory and due-process frameworks and the automated decision systems now operating within each sector's jurisdiction. This structural parallel, rather than being incidental, reflects a shared underlying challenge, algorithmic opacity and accountability, that neither sector's existing, independently developed governance apparatus was originally designed to address, suggesting that education and law's AI-governance challenges are more productively understood as a shared research problem than as two unrelated sector-specific concerns.

## FUTURE WORK

Future research should prioritize direct, comparative studies of algorithmic due-process and accountability mechanisms across education and law, testing whether governance frameworks developed for one sector, Citron's technological due process framework in particular, can be adapted to the other rather than each sector continuing to develop governance responses independently. Further work is needed to extend Selbst and Barocas's legal analysis of explanation adequacy specifically to educational AI applications, since the post-hoc explanation concerns documented for judicial risk-assessment algorithms apply with comparable force to AI-driven grading, admissions, and disciplinary decisions that current educational governance frameworks have not yet examined with similar rigor. Continued research should also pursue empirical evaluation of specific fairness-auditing methodologies across both sectors' consequential decision-support algorithms, extending the disparate-error-rate analysis Angwin, Larson, Mattu, and Kirchner conducted for COMPAS to analogous predictive systems in education, such as dropout-risk and admissions-prediction algorithms, where comparable demographic-disparity auditing remains considerably less developed. Finally, future work should examine whether a harmonized, cross-sector AI-governance standard, building on Floridi et al.'s cross-domain ethical framework, is a more effective policy path than the current pattern of sector-specific regulatory development, given this review's finding that education and law face substantially similar governance challenges despite operating under entirely separate regulatory traditions.

## REFERENCES

[1] O. Zawacki-Richter, V. I. Marín, M. Bond, and F. Gouverneur, "Systematic review of research on artificial intelligence applications in higher education—where are the educators?," International Journal of Educational Technology in Higher Education, vol. 16, art. 39, 2019.

[2] K. VanLehn, "The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems," Educational Psychologist, vol. 46, no. 4, pp. 197–221, 2011.

[3] G. Siemens, "Learning analytics: The emergence of a discipline," American Behavioral Scientist, vol. 57, no. 10, pp. 1380–1400, 2013.

[4] W. Holmes, M. Bialik, and C. Fadel, Artificial Intelligence in Education: Promises and Implications for Teaching and Learning. Boston, MA, USA: Center for Curriculum Redesign, 2019.

[5] B. P. Woolf, Building Intelligent Interactive Tutors: Student-Centered Strategies for Revolutionizing E-Learning. Burlington, MA, USA: Morgan Kaufmann, 2009.

[6] R. S. Baker and P. S. Inventado, "Educational data mining and learning analytics," in Learning Analytics: From Research to Practice, J. A. Larusson and B. White, Eds. New York, NY, USA: Springer, 2014, pp. 61–75.

[7] N. Selwyn, Should Robots Replace Teachers? AI and the Future of Education. Cambridge, U.K.: Polity Press, 2019.

[8] R. Susskind, Tomorrow's Lawyers: An Introduction to Your Future, 2nd ed. Oxford, U.K.: Oxford University Press, 2017.

[9] H. Surden, "Machine learning and law," Washington Law Review, vol. 89, no. 1, pp. 87–115, 2014.

[10] D. M. Katz, M. J. Bommarito II, and J. Blackman, "A general approach for predicting the behavior of the Supreme Court of the United States," PLOS ONE, vol. 12, no. 4, e0174698, 2017.

[11] D. K. Citron, "Technological due process," Washington University Law Review, vol. 85, no. 6, pp. 1249–1313, 2008.

[12] D. Remus and F. Levy, "Can robots be lawyers? Computers, lawyers, and the practice of law," Georgetown Journal of Legal Ethics, vol. 30, no. 3, pp. 501–558, 2017.

[13] J. Angwin, J. Larson, S. Mattu, and L. Kirchner, "Machine bias," ProPublica, May 23, 2016.

[14] R. M. Re and A. Solow-Niederman, "Developing artificially intelligent justice," Stanford Technology Law Review, vol. 22, no. 2, pp. 242–289, 2019.

[15] L. Floridi, J. Cowls, M. Beltrametti, R. Chatila, P. Chazerand, V. Dignum, C. Luetge, R. Madelin, U. Pagallo, F. Rossi, B. Schafer, P. Valcke, and E. Vayena, "AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations," Minds and Machines, vol. 28, no. 4, pp. 689–707, 2018.

[16] F. Pasquale, The Black Box Society: The Secret Algorithms That Control Money and Information. Cambridge, MA, USA: Harvard University Press, 2015.

[17] A. D. Selbst and S. Barocas, "The intuitive appeal of explainable machines," Fordham Law Review, vol. 87, no. 3, pp. 1085–1139, 2018.

[18] E. Zeide, "The structural consequences of big data-driven education," Big Data, vol. 5, no. 2, pp. 164–172, 2017.

[19] P. M. Regan and J. Jesse, "Ethical challenges of edtech, big data and student data: Protecting student privacy," Ethics and Information Technology, vol. 21, no. 3, pp. 167–179, 2019.

[20] K. D. Ashley, Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age. Cambridge, U.K.: Cambridge University Press, 2017.