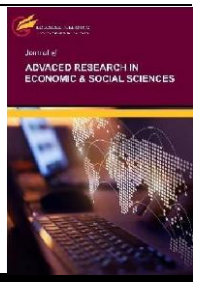
EDUGLOBE PUBLISHING
Expanding Knowledge Globally

Journal of Advanced Research in Economic & Social Sciences (JARESS)
Journal Homepage:www.eduglobepublishing.com



ROBOTS IN THE BACK OFFICE: A CRITICAL REVIEW OF ROBOTIC PROCESS AUTOMATION AND SERVICE ROBOTS IN EDUCATIONAL ADMINISTRATION AND MANAGEMENT

Ibrahim Oleolo[1*], Shakirat Okeowo[2], Alhussaini Wahab Salim Nasser[3].

1 Centre for Unmanned Technology International Islamic University Malaysia*

2 Univesiti Sultan Zainul Abideen, Terengannu Malaysia

3 College of Computing and Informatics Univesiti Tenegal Nasional

*Corresponding author: pdf29@iium.edu.my

Abstract: Educational robotics is among the best-reviewed fields in educational technology, yet almost the entire corpus concerns robots as objects of learning or agents of instruction. The parallel question of what happens when robots and robotic automation are deployed to run the institution rather than to teach in it has no consolidated literature. This critical review addresses that omission. Drawing on 60 sources spanning robotic process automation (RPA) in higher education and public administration, service and social robots in institutional non-teaching roles, and AI-enabled governance in educational leadership, it argues that a coherent research object exists, administrative robotisation in education, but that it is currently split across three literatures that do not cite one another. The review finds that the efficiency case rests on an evidentiary base that is structurally vendor-shaped; that gross hours released are reported while cost, fragility and failure are not; that the discretionary educational processes where automation matters most remain wholly unstudied; and that the administrative staff whose work is being automated are almost never the informants. It proposes a task-discretion-accountability framework for classifying which administrative processes should and should not be robotised, advances nine falsifiable propositions, and identifies K-12 district administration, cost-side accounting and longitudinal post-deployment evidence as the field's three most urgent priorities.

Key words: Educational management; robotic process automation; service robots; administrative automation; educational leadership; algorithmic discretion; critical review

1.0 Introduction

Two decades of research on robotics in education have produced a mature, well-mapped and repeatedly synthesised body of knowledge. Systematic reviews have established the effects of educational robotics on STEM competence, computational thinking, motivation and achievement; they have proposed frameworks for learner-robot relationships; and they have traced the field's evolution from robot construction toward broader robotics literacy. What every one of these reviews shares is a unit of analysis: the learner, and the robot's relationship to learning.

Educational institutions, however, are not only sites of instruction. They are organisations. They admit and enrol, timetable and allocate, procure and pay, record and report, comply and audit, discipline and safeguard. A very substantial share of the labour inside a school or university, and a very substantial share of the frustration expressed by those who work there, belongs to this administrative apparatus rather than to teaching. It is precisely this apparatus that is now the target of two converging technological pushes: robotic process automation (RPA), software bots that execute rule-based clerical transactions across existing systems; and physical service robots, embodied machines taking on reception, guidance, library operations, logistics, patrol and cleaning functions on campus.

The scholarly response to this convergence has been fragmented. Three literatures address parts of it, and they are almost hermetically sealed from one another. The first is an information-systems and computing literature on RPA, which contains real quantified claims about university back offices but is published in sustainability, information-systems and engineering

*Corresponding author: pdf29@iium.edu.my
https://eduglobepublishing.com/ All rights reserved.

venues, is exclusively concerned with higher education, and applies almost no critical governance or equity lens. The second is a human-robot interaction (HRI) and service-management literature on physical robots in institutional roles, which is engineering- and perception-shaped, measuring whether the machine works and whether bystanders like it, and which for educational settings must be reconstructed by analogy from hospitals, hotels and municipal citizen-service desks. The third is an educational leadership and management literature on AI in governance and decision-making, which possesses exactly the critical vocabulary the other two lack, including accountability, equity, transparency, distributed leadership and vendor influence, but which treats AI as software advice and contains essentially nothing on either physical robots or process automation. Where a robot appears in this literature at all, it appears as a survey prompt about robot teachers (Alkan et al., 2026), that is, as an instructional substitution question and not an administrative one.

The consequence is a structural blind spot. The efficiency claims generated in the first literature are never subjected to the governance critique developed in the third; the governance frameworks proliferating in the third are never tested against the concrete automated workflows described in the first; and the organisational realities documented in the second, that frontline staff often experience robots as additional work, that deployments depend on undocumented human rescue behaviour, and that accessibility is the weakest dimension of public-facing robots, reach neither.

1.1 Conceptual clarification: three layers of administrative robotisation

The term robot is doing a great deal of unacknowledged work in this domain, and much of the field's incoherence follows from its elasticity. This review distinguishes three layers, all routinely labelled robotics or automation in educational management discourse but differing in embodiment, in the kind of work they displace, and in the governance questions they raise.

Layer 1, software robots (RPA), comprises rule-based, non-embodied bots that emulate a human operator across existing user interfaces to execute clerical transactions: admissions data entry, enrolment processing, transcript generation, attendance recording, timetabling inputs, financial-aid processing, purchase orders, supplier setup, payroll and compliance checks (Bhardwaj & Kumar, 2025; Toma, 2023; Uta et al., 2026). What is displaced is clerical throughput. The governance question is what happens to the rule when the rule is executed by something that cannot notice that the rule does not fit.

Layer 2, embodied service robots, comprises physical machines performing institutional service functions: reception and greeting (Oruma, Colomo-Palacios, & Gkioulos, 2026), library shelf-reading and stock management (Tella & Ajani, 2022; Shahzad et al., 2025), campus security patrol (Ye & Robert, 2024; Marcu et al., 2023), logistics and internal delivery (Birkhoff et al., 2024), disinfection and cleaning (Fuszl et al., 2021), and citizen- or student-facing information provision (Baake et al., 2025). What is displaced, or more often rearranged, is embodied service labour. The governance question is who is accountable when the machine obstructs, fails, or excludes.

Layer 3, algorithmic decision support, comprises predictive and analytic systems informing institutional decisions: retention and at-risk modelling (Ayhan, 2026), resource allocation, admissions matching (Swist & Gulson, 2022), and the broader field of AI-enabled institutional governance (Jiang & Abdullah, 2026; Lipsou et al., 2026). What is displaced is judgement, and the governance question is the most consequential: whose discretion, exercised where, and reviewable by whom.

These layers are not merely a typology of convenience. They form an escalating sequence in which the object of automation moves from hands to bodies to minds, and in which the accountability problem sharpens at each step. The literature, revealingly, is thickest at Layer 3 in critical terms and thinnest there in empirical terms, while it is thickest at Layer 1 in efficiency claims and thinnest there in critique. This review's central diagnostic move is to read each layer with the other layers' strongest instruments.

2.0 Statement of the Problem

Educational institutions are procuring and deploying administrative automation now, on business cases built from figures whose provenance no one has audited, in a policy vacuum, and without an evidence base capable of telling a vice-chancellor or a district superintendent whether the decision is sound. The problem is not that the technology fails. Software bots do execute clerical transactions, disinfection robots do decontaminate more surfaces than manual cleaning, and chatbots do absorb routine enquiries. The problem is that the questions determining whether any of this is a good idea for a particular institution, what it costs in full, what happens when it breaks, who is accountable when it errs, whose burden it shifts, whose job it changes, and which decisions should never be handed to it, are precisely the questions the literature has not asked.

Three features of the field make this urgent. First, the two most-repeated efficiency figures in educational RPA scholarship originate in vendor marketing collateral and enter peer-reviewed synthesis without an independent verification step, so institutional business cases are being built on unaudited supplier claims. Second, the sharpest critical finding available, that rigid task delineation is misaligned with work requiring contextual professional judgement, has been established in social work but never tested on the discretionary educational processes where it matters most: admissions, extenuating circumstances, fee waivers, fitness to study and safeguarding referral. Third, the field is theoretically mis-scaled, applying individual-level acceptance models to what is characteristically an organisational procurement decision, and thereby producing a systematically depoliticised account of what happens when an institution automates.

Constituting administrative robotisation in education as a single research object, auditing the evidence advanced for it, and specifying what would have to be studied for the field to become useful to practising educational managers, is therefore the problem this review addresses.

2.1 Objectives of the Study

The objectives of this review are as follows.

i. Map how, and in what institutional roles, robotic process automation and physical service robots have actually been deployed in the administration of educational organisations.

ii. Appraise the quality, provenance and completeness of the evidence advanced for their organisational benefits.

iii. Examine how the literature treats the managerial dimensions of these deployments, including governance, procurement, accountability, workforce relations and administrative discretion.

iv. Evaluate the adequacy of the theoretical frameworks used to explain adoption and outcomes, and propose an alternative framework and research agenda.

2.2 Research Questions

The following research questions guided the review.

1. How, and in what institutional roles, have robotic process automation and physical service robots been deployed in the administration of educational organisations?

2. What is the quality, provenance and completeness of the evidence advanced for their organisational benefits?

3. How does the literature treat governance, procurement, accountability, workforce relations and the exercise of administrative discretion in these deployments?

4. Which theoretical frameworks are used to explain adoption and outcomes, and are they adequate to the organisational level at which the phenomenon operates?

3.0 Methodology

3.1 Review type and its justification

Grant and Booth's (2009) SALSA typology distinguishes fourteen review genres by their approach to Search, AppraisaL, Synthesis and Analysis. Three considerations rule out a conventional systematic review here. First, the phenomenon has no settled vocabulary: studies of the same practice are indexed variously as robotic process automation, intelligent process automation, administrative digitalisation, service robot deployment, AI-enabled governance and digital discretion, across education, information systems, public administration, HRI and service management, so that a protocol-driven search on any single controlled vocabulary would return a fraction of the relevant corpus and produce a spuriously precise flow diagram. Second, the corpus is heterogeneous by design, comprising PRISMA reviews, single-institution engineering proofs-of-concept, qualitative process studies, interview-based typologies, reliability calculations, conceptual frameworks, policy-document analyses and vendor case studies, with no common effect measure and nothing to pool. Third, and decisively, the review's purpose is appraisal and reconceptualisation rather than aggregation, which is the condition under which Grant and Booth locate the critical review and which Torraco (2005) describes for the integrative review: a genre whose output is a new conceptualisation of a body of work rather than a summary of it. Snyder's (2019) guidance on semi-systematic and integrative designs in management research provided the procedural discipline.

The review is therefore reported as a critical, configurative review, transparent about search and selection in the spirit of PRISMA 2020 (Page et al., 2021) without claiming its exhaustiveness. SPIDER (Cooke et al., 2012) rather than PICO structured the question, since the evidence of interest is organisational and largely qualitative. No bibliometric mapping layer was added, although Donthu et al. (2021) set out how one would be constructed; the corpus is too small and too dispersed across incommensurable indexing vocabularies for co-citation or co-word analysis to yield anything but an artefact of database coverage.

3.2 Search strategy

Searches were run across Scopus- and Web of Science-indexed content accessed via publisher platforms, Springer Nature Link, ScienceDirect, SAGE Journals, Emerald Insight, Taylor and Francis Online, Wiley Online Library, MDPI, Frontiers, the ACM Digital Library, IEEE Xplore, ERIC, DOAJ and Google Scholar, supplemented by citation chaining in both directions and by targeted retrieval of vendor case material. Four search clusters were combined with Boolean operators and truncation. Cluster A (software automation) combined “robotic process automation” OR RPA OR “intelligent process automation” with education OR universit* OR school OR “higher education” OR administration OR “back office” OR admissions OR enrolment. Cluster B (embodied robots) combined “service robot” OR “social robot” OR “humanoid robot” OR “security robot” OR “delivery robot” OR “disinfection robot” with campus OR universit* OR library OR school OR reception OR “front desk” OR institution*. Cluster C (management and governance) combined “educational administration” OR “school leadership” OR “educational management” OR “higher education governance” with automation OR “artificial intelligence” OR algorithm* OR robot*. Cluster D (theory and discretion) combined “street-level bureaucracy” OR “digital discretion” OR “artificial discretion” OR “administrative burden” with automation OR ICT OR “artificial intelligence”.

3.3 Inclusion, appraisal and corpus

Sources were included where they addressed a non-instructional, organisational deployment of robotic automation or embodied robots, or where they supplied theoretical or methodological apparatus for analysing such deployments. Peer-reviewed work published between 2017 and 2026 was prioritised, with seminal earlier work retained where genealogically necessary (Davis, 1989; DiMaggio & Powell, 1983; Meyer & Rowan, 1977; Bovens & Zouridis, 2002; Lipsky, 2010; Rogers, 2003; Tornatzky et al., 1990). Analogue-sector evidence from healthcare, hospitality, libraries and municipal government was included where the educational evidence was absent and the organisational structure comparable, and is labelled as analogue throughout.

Vendor and practitioner material was deliberately retained rather than excluded, but is quarantined into a separate evidential category and is never cited in support of an empirical claim. This is a methodological decision with substantive consequences: excluding grey literature would have concealed the review's single most important finding about the provenance of the field's headline numbers. Appraisal was conducted narratively against four criteria: whether an empirical claim is traceable to primary measurement or to a secondary report; whether a baseline and a measurement method are stated; whether the affected workforce is among the informants; and whether cost, failure and discontinuation are addressed alongside benefit. Sixty sources are cited, comprising 21 on RPA and administrative automation, 21 on embodied service robots and workforce relations, 16 on AI and governance in educational leadership, and the remainder supplying theoretical and methodological anchors. Two vendor case studies are cited as vendor claims only.

3.4 Theoretical lens

The overwhelming majority of adoption studies in this space are built on the Technology Acceptance Model (Davis, 1989) and its unified successors (Venkatesh et al., 2003, 2012), sometimes re-specified for AI agents as the AI Device Use Acceptance model (Gursoy et al., 2019) or for embodied service encounters as the Service Robot Acceptance Model (Wirtz et al., 2018). Bhardwaj and Kumar (2025), for instance, report that adoption success in higher-education RPA hinged on perceived usefulness, ease of use and staff training: TAM constructs, essentially unmodified, applied to an institutional programme. This is a mis-scaling, and the point has been made by the framework's own originators. Venkatesh et al. (2016) concede the insufficiency of single-level individual-intention models and propose a multi-level alternative; Bagozzi (2007) had already argued that TAM is deterministic and under-theorises group and organisational decision processes; and O'Dea et al. (2024), writing for an education audience, restate the two enduring problems of over-reliance on self-reported behavioural intention and loss of parsimony through indefinite model extension. The mis-scaling matters more here than in most domains, because the decision to robotise an administrative process is characteristically not an individual user's decision at all. It is a procurement decision, taken by a senior team, shaped by vendor relationships, sector benchmarking and audit expectation, and imposed on the frontline staff whose acceptance is subsequently measured.

This review therefore reads the corpus through a composite lens drawn from four alternative anchors. The Technology-Organisation-Environment framework (Tornatzky et al., 1990; Oliveira & Martins, 2011) locates adoption in the interaction of technological, organisational and environmental conditions; its value is demonstrated by Huynh and Lindgren (2026), whose TOE-coded review of 34 RPA studies attacks the vendor narrative of RPA as a lightweight technology and derives twelve challenges and seven sequential preconditions. Sociotechnical systems and affordance theory (Sarker et al., 2019; Leonardi, 2011; Volkoff & Strong, 2013) supplies what acceptance models cannot, namely an explanation of variation in outcomes across institutions with similar technology; the most theoretically sophisticated study in the education-specific corpus, Doolin et al.'s (2025) process study of a university RPA initiative, works in exactly this register. Institutional theory (DiMaggio & Powell, 1983; Meyer & Rowan, 1977) explains why institutions adopt the same platforms irrespective of demonstrated fit and why announced policies may bear little relation to operational practice; Rughinis et al. (2025) and Khantong and Sriboonlue (2026) both document this mechanism without naming it. Finally, the system-level bureaucracy tradition, the anchor the educational management literature most conspicuously lacks, establishes that frontline discretion constitutes de facto policymaking (Lipsky, 2010), that information technology migrates that discretion to system designers with consequences for constitutional control (Bovens & Zouridis, 2002), that digital discretion strengthens ethical and democratic values while weakening professional and relational ones (Busch & Henriksen, 2018), and that artificial discretion can be classified by task complexity against equity and manageability costs (Young et al., 2019). Herd and Moynihan (2018) add the distributive dimension, and Dunleavy et al. (2006) the macro-frame of digital-era governance.

4.0 Results: Thematic Synthesis

4.1 The efficiency claim and the laundering of vendor evidence

The strongest and most frequently repeated claim in this field is that administrative robotisation delivers large, quantified efficiency gains. Bhardwaj and Kumar's (2025) PRISMA review of 54 studies reports time reductions up to 99.9% on internship reporting and attendance checking, cost savings of 43 to 67%, and institutional labour savings of roughly 10,000 hours per year at the University of Melbourne and 23,000 hours per year at the University of Auckland. These are the numbers that circulate: they appear in subsequent reviews, in practitioner discourse, and in the business cases institutions write for their own boards.

A provenance audit of these figures is sobering. The Auckland and Melbourne figures originate, respectively, in UiPath and Automation Anywhere customer-story pages (UiPath, n.d.; Automation Anywhere, n.d.). These are marketing collateral. They state no baseline methodology, no measurement period, no monetary accounting, no implementation cost and no independent audit. The Auckland page additionally claims 96% email-routing accuracy and a 99% finance-process success rate; the Melbourne page claims 97% throughput on supplier-detail processing across 22 automated processes. Every one of

these numbers is self-reported by the party selling the software and the party that bought it. The 99.9% time-reduction figure traces to two conference papers indexed in IEEE Xplore on internship reporting and attendance checking, neither of which could be retrieved in full for this review: the single most spectacular statistic in the field rests on documents that a determined reviewer could not read.

What has occurred is a laundering cycle. A vendor publishes an unaudited figure; a peer-reviewed review includes it as a documented case; subsequent work cites the peer-reviewed review; and the figure acquires the epistemic status of a research finding while never having acquired the properties of one. This is not a criticism of Bhardwaj and Kumar's diligence, since their review is a legitimate synthesis of what has been published, but of a field whose published record is thin enough that vendor cases are load-bearing. The pattern is not new and was diagnosed at the origin of the field: Syed et al. (2020), in the agenda-setting review of RPA scholarship, explicitly noted the early evidence base's heavy reliance on vendor white papers. The intellectual lineage compounds the problem, since the framing of RPA as a strategic transformation lever, later imported wholesale into university shared-services discourse, derives from Willcocks et al.'s (2017) teaching case on a private-sector insurance back office, a pedagogical text produced within an industry-sponsored research stream.

Set against this, the more rigorously bounded empirical work reports far more modest and more honest figures. Uta et al. (2026), in a mixed-methods case study of a multilingual engineering faculty combining VR, chatbot and RPA components, report a chatbot logging 1,283 interactions over seven months, offsetting an estimated 107 administrative hours, and an RPA routine reducing password-reset handling from several minutes to under one minute with fewer errors. That is a real, plausible, verifiable gain, and it is three orders of magnitude away from the headline claims. Marchena Sekli and Godo's (2025) PRISMA review, screening 229 records down to 17 empirical studies, explicitly notes the small evidence base, the dominance of single-institution case studies, the near-absence of outcome evidence, and unexamined scalability and cost-effectiveness. The field's efficiency literature thus exhibits a systematic asymmetry: the most-cited numbers have the weakest provenance, and the best-designed studies report the smallest effects.

4.2 The missing denominator: cost, fragility and failure

Across the corpus, no study reports total cost of ownership for an educational RPA deployment, covering licensing, implementation consultancy, process re-engineering, bot maintenance and the recurring cost of repair when an upstream system changes its interface. No study reports bot breakage rates. No study reports an abandoned, reversed or discontinued deployment, despite the near-certainty that such cases exist. And no study establishes what the released hours were converted into, whether reduced headcount, redeployment to higher-value work, or simple absorption into an unchanged workload.

The analogue literatures make the omission conspicuous by contrast, because they measure exactly what education does not. Birkhoff et al. (2024) observed two logistics cobots complete 89 deliveries over seven weeks in a 900-bed hospital and counted the human interventions required to make that happen: 107 interventions, 72 of them simply to operate elevators. They concluded that cobots remain in an early innovation phase and require human support to function, and, tellingly, explicitly did not measure workflow productivity or construct a business case. Fuszl et al. (2021), field-testing a UV-C disinfection robot in routine hospital cleaning, found it decontaminated 96.9% of surfaces against 50.0% for manual cleaning alone, yet judged it a delicate high-tech device not yet ready for autonomous use in clinical routine, because its interface demanded advanced technical knowledge and it frequently required operator intervention. Misaros et al. (2024), calculating the reliability of a Pepper humanoid handling office documents, report a headline system reliability of approximately 0.99993 over 10,000 hours, but a component-level mean time between failures of 1,667 hours for the manipulation subsystem and 22.62 hours for the safety subsystem. A safety subsystem that fails roughly once a day is not a footnote; in an institution owing a duty of care to minors or to disabled users, it is the whole story.

Two further findings deserve prominence. Kraus et al. (2024) derive a five-part typology of RPA implementation challenges from consultant interviews, spanning process-related, technical, resource-related, psychological and coordinative categories. Khantong and Sriboonlue's (2026) PRISMA review of 83 articles finds that the psychological and coordinative barriers, including status quo bias and dehumanisation concerns, are the most systematically underestimated failure modes, ranking above technical ones, and that public-sector contexts require institutional-legitimacy framing rather than cost-reduction framing. Lindgren et al. (2022), studying three Swedish municipalities, identify top-down implementation without grassroots buy-in, weak process-design and IT-infrastructure competence, over-dependence on individual champions and external consultants, and difficulty identifying suitable processes, concluding that the binding constraint is structural misalignment between organisational sub-units rather than technology. Juell-Skielse et al. (2022), surveying Swedish public organisations nationally, find high awareness of RPA alongside modest actual implementation and a marked adoption gap between central and local government. Durao and dos Reis (2024) map perceived value across the value chain from seven expert interviews, finding credible internal and downstream value but inconclusive upstream effects.

A literature that reports benefits without costs, gross savings without net savings, and successes without failures is not an evidence base; it is a promotional genre with citations. The absence of a single documented discontinuation in the educational corpus should be read not as evidence that none has occurred but as evidence of what the field considers publishable. Educational management research should recognise this pattern, since it is the same publication asymmetry that once distorted school improvement and educational technology effectiveness research, and the same corrective applies: cost-effectiveness accounting, pre-registered evaluation, and publication of null and negative results.

4.3 Discretion, equity and the processes nobody has studied

The most consequential critical finding available to this review is not from education at all. Germundsson et al. (2024), analysing group interviews across four Swedish municipalities implementing RPA in social assistance, found that RPA does not inherently lead to the expected all-encompassing enhancements of faster and fairer eligibility determinations, and that the rigid task delineation RPA requires is fundamentally misaligned with casework depending on individualised professional judgement and contextual knowledge of the client.

Transposed to education, that finding's force is immediate, because educational administration is saturated with exactly this kind of work. In admissions and school choice, Swist and Gulson (2022) show that the New York City matching algorithm operates within uncertain archives that embed existing inequalities through classification choices, masking spatial and socioeconomic disparity while distributing a veneer of choice, and that youth activists surfaced residual categories the algorithm excluded, such as census tract and free-lunch eligibility. Extenuating circumstances, mitigating circumstances and fitness-to-study decisions exist precisely because rules do not fit every case. In fee waivers, hardship funds and financial-aid discretion, Herd and Moynihan's (2018) analysis predicts that automating compliance requirements will redistribute learning and psychological costs onto the least-resourced applicants. In attendance enforcement and safeguarding referral, a rule-based trigger and a professional judgement can diverge with serious consequences. And in retention and at-risk intervention, Ayhan (2026), modelling 523 students across three universities, reports a best-performing multilayer perceptron reaching 70% accuracy and, crucially, notes that the model's errors differed systematically from instructor judgements. A 70%-accurate model whose errors are non-randomly distributed relative to human judgement is not a neutral efficiency gain; it is a reallocation of who gets misjudged.

No study in the reviewed corpus tests automation against any of these processes. This is the field's single largest substantive gap, and the theoretical apparatus for analysing it already exists and is simply not being used. Wang (2024) puts the governance stake sharply: algorithmic decisions in education are policy decisions, implemented by private vendors with significant public impact and minimal regulatory oversight, and made ungovernable by three properties, namely speed of development, lack of interpretability, and unpredictable emergence of new capabilities. Jiang and Abdullah's (2026) PRISMA review of 27 governance studies confirms the imbalance empirically: decision-support systems and predictive analytics dominate, operational efficiency and predictive accuracy are the most emphasised outcomes, while transparency, accountability, trust, equity and workload redistribution remain comparatively neglected.

The readiness of the people expected to operate under these arrangements is, on the available evidence, low. Alsuhaymi and Eltantawy (2026), surveying 173 principals, supervisors and teachers across Riyadh public special education schools, found high agreement on AI's administrative importance ($M = 3.96$) coexisting with markedly low confidence on ethical considerations ($M = 2.38$) and only moderate acceptance of AI in a decision-support role ($M = 3.07$), which respondents regarded as supplementary to professional judgement; training exposure mattered substantially, while professional role and years of experience did not. Alkan et al. (2026), surveying 324 Turkish school principals, similarly found that a majority lacked adequate awareness of big data and data mining, with educational level the strongest predictor of technological literacy. Zeng et al. (2025), reviewing 14 empirical studies, find that principals shape adoption chiefly through resource allocation, professional development and learning-community building, with combined transformational and digital leadership most effective, but the review cannot establish causal links.

The educational leadership literature has developed a rich normative vocabulary for these problems. Lipsou et al. (2026) recommend embedding AI in distributed leadership, mandating explainability and audits in procurement, and investing in privacy and data literacy; DeMatthews et al. (2026) offer six AI leadership domains aligned to professional standards; Polat et al. (2025) map ethical AI themes bibliometrically; Arar et al. (2025) synthesise the field integratively and Arar et al. (2026) map it bibliometrically. But it applies this vocabulary almost exclusively to teaching-adjacent AI, and it applies it prescriptively. Frameworks proliferate, with Marquez et al. (2024) counting fourteen critical success factors for learning analytics adoption and concluding that no single framework covers all dimensions, without anyone testing whether adopting one changes anything. A field that has produced more bibliometric maps of itself than empirical studies of the phenomenon it maps has a sequencing problem.

4.4 The absent workforce

Administrative and professional-services staff are the people whose work is being automated, and they are almost never the informants. Kraus et al. (2024) interviewed consultants; Durao and dos Reis (2024) interviewed seven RPA experts; Guarda et al. (2025) review the literature conceptually with no primary data. Doolin et al. (2025) stand nearly alone in studying an actual university workforce in situ, theorising implementation as socio-technical change across five patterns of initiation, mobilisation, configuration, adaptation and evaluation.

The analogue literature is more candid, and its findings should be uncomfortable reading for anyone building a business case. Macalupu et al. (2025), studying staff and service robots in aged care, document the sharpest management-relations finding in the corpus: industry leaders regarded robots as a long-term positive investment while frontline staff experienced them as disruptive to established routines, reporting frequent malfunctions, inadequate training, and increased workload from managing the robots rather than benefit from them; staff also objected to abrupt robot speech such as “Make way, robot coming.” Paluch et al. (2021), from 36 interviews, construct a willingness-to-collaborate framework and four employee personas, supporter, embracer, resister and saboteur, thereby naming active staff resistance as a predictable deployment outcome rather than an anomaly. Horpynich et al. (2025), in an SEM analysis of 279 hospitality workers, find that negative service-robot awareness and perceived risk significantly raise turnover intentions, mediated by job insecurity, with marked

generational differences. Shahzad et al. (2025), reviewing robot adoption in academic libraries, list fear of job loss and shortage of skilled manpower among the principal barriers alongside financial hurdles and absent organisational planning.

Against the alarmist reading, Guarascio et al. (2025), meta-analysing 33 employment and 19 wage studies, find aggregate employment and wage effects negligible, with negative effects concentrated in US manufacturing among middle-skilled workers, and detect publication bias favouring negative wage findings. But this evidence rests on industrial-robot exposure measures in manufacturing economies and cannot be transferred to campus reception, library or cleaning staff, a caution the education literature would need to observe before either reassuring or alarming its workforce.

Two sources point toward what a constructive alternative would look like. Ohneberg et al. (2025), interviewing nursing professionals about where robots might fit, found staff willing and confident about operating robotic systems provided training and ongoing support exist, and identified implementation preconditions spanning spatial infrastructure, person-centred factors, organisational readiness and technology, a preconditions list generated by the workers rather than by the vendor. Oladokun et al. (2026) address the skills question directly for librarians working alongside robotic systems. Both suggest that the determinant is not staff attitude but institutional preparation, which is precisely the variable the education literature does not measure. Nowhere in the corpus is there a study reporting union consultation, job redesign, redeployment, renegotiated job descriptions, or a collective agreement following a robot deployment in an educational institution. Ozturkcan (2025) argues conceptually that robots are not neutral actors but can reinforce existing biases through design, deployment and interaction scripts, with outcomes mediated by HR policy, leadership, training and sociotechnical adaptation, and notes that longitudinal and intersectional empirical work does not exist. Freeing staff time is thus the field's central normative claim and its least examined one: asserted by managers and consultants, measured as an aggregate, and never traced to the individuals whose time was freed.

4.5 The governance and procurement vacuum

If the workforce is absent from the literature, the managerial apparatus around the machine is very nearly invisible. No study in this corpus follows a procurement decision from business case through tender to contract. None reports what a campus reception or patrol robot cost against the staff time it displaced. None examines vendor lock-in, maintenance contracts, or the service-level agreements that in practice determine whether a deployment survives past pilot. Governance appears only in the abstract: Oruma, Ayele, Sechi and Rodsethol (2023), mapping 40 studies on security aspects of social robots in public spaces, find a substantive gap between industrial robot safety standards and what institutionally deployed robots actually require, with cybersecurity dominating 35% of studies, ethics 22.5% and legal frameworks 17.5%, and propose a combined framework of physical safety, communication encryption, authentication, operational transparency and social-robot-specific ethical guidance. But mapping the standards that ought to exist is not the same as observing an institution operating a policy, running an ethics review, completing a data-protection impact assessment, or maintaining an incident-reporting route for a deployed robot.

Accountability, when failure occurs, is nowhere assigned to a named institutional role. Consider what the corpus documents without allocating responsibility: 107 human interventions to keep two hospital cobots moving (Birkhoff et al., 2024); a safety subsystem with a 22.62-hour mean time between failures (Misaros et al., 2024); a university reception robot scoring 89.7% on physical-safety confidence and 90.0% on data protection but only 60.6% on accessibility, with explicit barriers for wheelchair users and those needing alternative communication modalities (Oruma, Colomo-Palacios, & Gkioulos, 2026); and a sidewalk delivery pilot in which robots obstructed pavements, distracted children and dogs, forced pedestrians to reroute, and required bystander rescue when displaying “HELP ME,” while limited public awareness of the pilot allowed misinformation to spread and disabled pedestrians were underrepresented in the study despite facing disproportionate obstruction risk (Weinberg et al., 2023). For an educational institution, that accessibility finding is not a usability detail but a potential breach of statutory duty: a reception robot that works less well for disabled students is a differential service, and the literature contains no account of any institution having assessed one on that basis.

There is also almost no evidence that institutions set success criteria in advance, and the studies inherit this, measuring microbial counts, engagement rates, reliability coefficients and attitude scales rather than whether an organisational objective was met. Baake et al. (2025) provide the rare and instructive exception, a head-to-head comparison of a Pepper robot against the cheaper incumbent technology, a conventional information terminal, across 37 participants in Lubeck public institutions: the robot scored higher on hedonic qualities, but the terminal was rated more pragmatic and was preferred overall. That a robot loses to a touchscreen when someone bothers to run the comparison is precisely the kind of finding the field needs more of, and it is nearly unique in the corpus. Even the field's in-the-wild evidence is thinner than the phrase implies: Arreghini et al. (2024) deployed a service robot across two crowded public locations with 1,777 participants and showed that proactive, behavioural-cue-triggered engagement substantially outperforms passive or distance-triggered modes, a genuinely useful design finding generated over three days totalling just over five hours of run time, with the robot offering chocolate rather than performing any institutional service function. Vendor dependency, meanwhile, is flagged repeatedly and studied never: Lipsou et al. (2026) and Gorman and Brennan (2026) both name vendor influence as a leadership concern, and Rughinis et al. (2025) theorise dual black-boxing, algorithmic opacity compounded by usage invisibility, yet no one has examined an actual education technology procurement.

4.6 Structural gaps in the evidence base

Four structural features cut across the themes above. The first is the K-12 void. Every credible peer-reviewed item located in this review sits in higher education or in municipal government, while school-level and district-level administration, the

sector employing the largest number of educational administrators worldwide, is empirically empty. Claims about reducing teacher administrative burden through automation rest on consultancy material, not research. This is doubly striking because the discretionary, safeguarding-sensitive processes identified in Section 4.3 are most concentrated in compulsory schooling.

The second is that two literatures do not speak. The RPA literature and the educational leadership literature do not cite each other: efficiency claims are never subjected to governance critique, and governance frameworks are never tested against concrete automated workflows. Constituting the shared object is a precondition for either literature improving. The third is geographic and methodological narrowness. Empirical work clusters in Sweden, Australasia, Romania, Turkey, Saudi Arabia and Ireland; designs are overwhelmingly single-case and cross-sectional; there is no longitudinal, comparative or quasi-experimental study of administrative robotisation in education, no cost-effectiveness analysis, and no multi-institution study of the same deployment. Several reviews exclude ERIC and all non-English sources. The Global South is essentially unrepresented, a serious omission given that resource constraint, infrastructure variability and administrative burden are most acute there and that isomorphic pressure to adopt Northern platforms is strongest.

The fourth is technological displacement of the object itself. Both Huynh and Lindgren (2026) and Khantong and Sriboonlue (2026) note that rule-based RPA is being absorbed into LLM-driven and agentic automation, yet no education-sector work addresses what happens to existing rule-based bot estates under that transition, how they are governed, or whether they remain auditable. A field that has not yet produced adequate evidence about the technology it is studying now faces that technology being superseded. Finally, the higher-education robotics literature is so dominated by teaching and learning that Budagov et al. (2024), reviewing 30 articles on robot assistants in higher education, find telepresence, writing-course and tutoring applications and administrative deployment barely at all. That the administrative story on campus has to be reconstructed by analogy from hospitals, hotels and citizen-service desks is itself the review's most compact finding.

5.0 Discussion

5.1 A task-discretion-accountability framework

The synthesis above suggests that the field's central practical question has been posed wrongly. It has asked whether administrative processes in education can be robotised, a question to which the answer is trivially yes for a large class of tasks. The question worth asking is which processes should be, under what governance, and with what recourse when the automation is wrong. The framework proposed here classifies administrative processes along three dimensions drawn from the theoretical anchors set out in Section 3.4.

The first dimension is task determinacy (after Young et al., 2019, and Bovens & Zouridis, 2002): is the process governed by a complete rule that a system can execute without needing to recognise an unanticipated case? Password resets, transcript generation, supplier setup and purchase-order matching are determinate; extenuating circumstances, safeguarding referral and hardship awards are not. The system-level bureaucracy thesis predicts that automating an indeterminate process does not eliminate discretion but relocates it into the specification, where it becomes invisible and unreviewable. The second dimension is consequence asymmetry (after Herd & Moynihan, 2018): what is the distribution of harm when the process errs? A misrouted invoice is recoverable and its cost falls on the institution, whereas a wrongly refused fee waiver, a missed safeguarding trigger, or a retention model whose errors are systematically distributed (Ayhan, 2026) imposes non-recoverable harm on individuals, concentrated among the least-resourced. Consequence asymmetry, not process volume, should determine the governance tier. The third dimension is accountability locatability: can a named institutional role be held answerable for an automated output, with a documented route to human review? Where the answer is no, as it is throughout the reviewed corpus, the process should not be automated regardless of how determinate it is.

Table 1

Task-Discretion-Accountability Framework for Administrative Robotisation in Education

	Low consequence asymmetry	**High consequence asymmetry**
High determinacy	Automate. Standard change control. Measure net cost, not gross hours. Examples: transcript generation, supplier setup, password resets, attendance recording.	Automate with audit. Mandatory logging, periodic bias audit, documented human-review route, disability-impact assessment. Examples: fee calculation, enrolment eligibility, compliance reporting.
Low determinacy	Augment, do not automate. Decision support with the professional retaining the decision. Examples: timetable optimisation, resource allocation modelling.	Do not automate. Human decision, with automation confined to information assembly. Examples: extenuating circumstances, fitness to study, safeguarding referral, hardship awards, exclusions.

The framework's contribution is to make the default explicit. The field's current default, which is to automate wherever technically feasible and evaluate afterwards by counting hours, inverts the appropriate ordering. Under the framework, the burden of proof sits with the proposer of automation in the high-asymmetry rows, and gross-hours metrics are inadmissible as evidence anywhere.

5.2 A research agenda

The following nine propositions are stated so as to be falsifiable and are ordered by urgency.

1. Provenance. A citation-tracing study of efficiency claims in educational RPA literature will find that a majority of quantified benefit figures resolve, within three citation steps, to vendor-authored material with no stated baseline methodology.
2. Net effect. Where total cost of ownership is measured, covering licensing, consultancy, re-engineering, maintenance and breakage, net savings from educational RPA deployments will be substantially smaller than reported gross hours released, and a non-trivial minority of deployments will show negative net returns.
3. Conversion. In most deployments, hours released will not convert into headcount reduction or measurable redeployment to higher-value work, but will be absorbed by unchanged or expanded workload.
4. Discretion. In educational processes of low determinacy, automation will relocate rather than eliminate discretion, shifting it to specification writers and vendors, and will reduce the reviewability of individual decisions.
5. Distribution. Automation of eligibility and compliance processes will redistribute administrative burden toward applicants with least capacity to absorb learning and psychological costs, widening rather than narrowing equity gaps.
6. Workforce. Frontline administrative staff will report increased rather than decreased task load in the first twelve months of embodied-robot deployment, with the discrepancy between leadership and frontline assessments largest where staff were not involved in specification.
7. Isomorphism. Adoption decisions across comparable institutions will be better predicted by peer adoption, sector benchmarking and vendor relationships than by institution-specific process analysis or documented need.
8. Decoupling. Institutions with published AI and automation governance policies will not differ systematically from those without in observable operational practice, such as ethics review, impact assessment and incident logging, absent an external audit requirement.
9. The incumbent comparison. Where embodied service robots are compared head-to-head against cheaper incumbent technologies on pragmatic service quality, the incumbent will match or outperform the robot in a majority of institutional service contexts.

The agenda requires, minimally, multi-site longitudinal case studies of educational RPA programmes running at least 24 months and including at least one discontinued deployment; a cost-effectiveness study using a stated total-cost-of-ownership protocol; K-12 district-level empirical work, which currently does not exist; studies whose informants are administrative and professional-services staff rather than leaders, consultants or vendors; disability-impact and accessibility evaluation of any student-facing embodied deployment, given the 60.6% accessibility confidence reported by Oruma, Colomo-Palacios and Gkioulos (2026); and Global South evidence, without which the field's claims are parochial.

5.3 Limitations

This review is bounded in four ways. It is English-language only, which given the concentration of empirical work in Sweden, Turkey, Romania and Saudi Arabia almost certainly excludes relevant non-English scholarship. Several major publisher platforms intermittently refused programmatic access during searching, so a minority of bibliographic details were verified through repositories and indexing services rather than publisher records, and sources whose details could not be fully resolved were demoted from evidential citation to discussion. As a critical rather than systematic review it makes no exhaustiveness claim and reports no inter-rater reliability, since screening and appraisal were narrative and single-reviewer, a limitation shared with several of the reviews it critiques, notably Jiang and Abdullah's (2026) acknowledged single-coder screening, and one that a subsequent protocol-driven review should correct. Finally, the reliance on analogue-sector evidence is a substantive interpretive risk: educational institutions differ in their duty of care, their public accountability and their professional cultures, and findings transposed from a hotel or a hospital ward should be treated as hypotheses for educational testing rather than as established results.

6.0 Conclusion

The literature on robotics in education has, for two decades, asked what robots do to learners. It has scarcely begun to ask what robots do to the organisations that educate them. This review has argued that a coherent research object, administrative robotisation in education, exists, is being actively deployed in institutions, and is currently theorised by nobody, because the three literatures with a claim on it do not read each other.

The picture that emerges is not one of a technology that does not work. Software bots do execute clerical transactions, disinfection robots do decontaminate more surfaces than manual cleaning, and chatbots do absorb routine enquiries. It is a picture of a field that has not yet learned to ask the questions that determine whether any of this is a good idea for a particular institution: what it costs in full, what happens when it breaks, who is accountable when it errs, whose burden it shifts, whose job it changes, and which decisions should never be handed to it at all.

Those are, recognisably, educational management questions. They are the questions the field already knows how to ask about accountability regimes, about workload, about equity of provision, and about the gap between policy text and institutional practice. The most useful thing educational management scholarship can do here is not to develop yet another adoption framework, since there are already more frameworks than there are empirical studies, but to bring its existing critical instruments to bear on a phenomenon that has so far been evaluated almost entirely by the people selling it.

7.0 Recommendations

Based on the findings of this review, the following recommendations are made.

1. Institutional leaders should treat gross-hours-saved figures as unverified vendor claims until a baseline methodology is produced, and should require any automation business case to state total cost of ownership, maintenance liability, the named role accountable for automated outputs, and the human-review route.
2. Institutions should apply the task-discretion-accountability framework before automating, recognising that the operative question is not technical feasibility but determinacy and consequence asymmetry; where a process involves discretion exercised on behalf of a student in difficulty, automation should assemble information rather than decide.
3. Procurement offices should specify audit rights, explainability obligations, exit and data-portability terms, and accessibility conformance before signature, since procurement is the single moment at which the institution holds contractual power and the literature offers no example of an institution having used it.
4. Management should involve the administrative and professional-services staff whose processes are being automated in specification, and should evaluate deployments partly on those staff's reported task load, treating resistance as a predictable response to imposed change rather than as a pathology.
5. Ministries and system-level authorities should set floor requirements for consequential automated decisions in education, including mandatory logging, a statutory right to human review, accessibility conformance for public-facing embodied systems, and incident reporting, with a standards conversation for robots deployed in settings involving children.
6. Educational management journals should commission empirical work on administrative robotisation, and reviewers should apply to automation claims the same evidential standards the field already applies to school effectiveness claims.

References

Alkan, B. B., Inal, G., Karakus, L., & Alkan, N. (2026). Shaping the future of education: School principals' views on AI, big data and robot teachers. AI & Society, 41(2), 1417-1433. https://doi.org/10.1007/s00146-025-02570-w

Alsuhaymi, A., & Eltantawy, M. M. (2026). Utilization of artificial intelligence to support administrative decision-making in special education institutions in Saudi Arabia: Perceptions of principals, supervisors, and teachers. Frontiers in Artificial Intelligence, 9, Article 1825189. https://doi.org/10.3389/frai.2026.1825189

Arar, K., Tlili, A., & Salha, S. (2026). Human-machine symbiosis in educational leadership in the era of artificial intelligence (AI): Where are we heading? Educational Management Administration & Leadership, 54(3), 669-692. https://doi.org/10.1177/17411432241292295

Arar, K., Tlili, A., Schunka, L., Salha, S., & Saiti, A. (2025). Reimagining educational leadership and management through artificial intelligence: An integrative systematic review. Leadership and Policy in Schools. https://doi.org/10.1080/15700763.2025.2451982

Arreghini, S., Abbate, G., Giusti, A., & Paolillo, A. (2024). A service robot in the wild: Analysis of users intentions, robot behaviors, and their impact on the interaction (arXiv:2410.03287). arXiv. https://doi.org/10.48550/arXiv.2410.03287

Automation Anywhere. (n.d.). University of Melbourne RPA case study [Vendor case study]. https://www.automationanywhere.com/resources/customer-stories/university-of-melbourne

Ayhan, I. (2026). Strategic use of predictive analytics for student retention in blended higher education: Implications for education management. Journal of Higher Education Policy and Management. https://doi.org/10.1080/1360080X.2026.2643372

Baake, E., Konig, F., Pott, C., Sievers, T., & Heine, M. (2025). User experience comparison of a social robot vs. info terminal in citizen services. In Social Robotics (ICSR + AI 2024) (pp. 12-23). Springer. https://doi.org/10.1007/978-981-96-3525-2_2

Bagozzi, R. P. (2007). The legacy of the technology acceptance model and a proposal for a paradigm shift. Journal of the Association for Information Systems, 8(4), Article 3, 244-254.

Bhardwaj, V., & Kumar, M. (2025). Transforming higher education with robotic process automation: Enhancing efficiency, innovation, and student-centered learning. Discover Sustainability, 6(1), Article 356. https://doi.org/10.1007/s43621-025-01198-6

Birkhoff, S. S., Merring, P., Spence, A., Bassett, W., & Roth, S. C. (2024). Integrating collaborative robots into a complex hospital setting: A qualitative descriptive study. Delaware Journal of Public Health, 10(5), 20-27. https://doi.org/10.32481/djph.2024.12.05

Bovens, M., & Zouridis, S. (2002). From street-level to system-level bureaucracies: How information and communication technology is transforming administrative discretion and constitutional control. Public Administration Review, 62(2), 174-184. https://doi.org/10.1111/0033-3352.00168

Budagov, F., Leoste, J., Meeran, M. T., & Robal, T. (2024). A systematic literature review on applicability of robot assistants in higher education. In Methodologies and Intelligent Systems for Technology Enhanced Learning (MIS4TEL 2024) (pp. 21-32). Springer. https://doi.org/10.1007/978-3-031-73538-7_3

Busch, P. A., & Henriksen, H. Z. (2018). Digital discretion: A systematic literature review of ICT and street-level discretion. Information Polity, 23(1), 3-28. https://doi.org/10.3233/IP-170050

Cooke, A., Smith, D., & Booth, A. (2012). Beyond PICO: The SPIDER tool for qualitative evidence synthesis. Qualitative Health Research, 22(10), 1435-1443. https://doi.org/10.1177/1049732312452938

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. MIS Quarterly, 13(3), 319-340. https://doi.org/10.2307/249008

DeMatthews, D., Reyes, P., Hart, T. D., & James, L., III. (2026). Leadership for artificial intelligence use in schools: A six-domain framework for ethical, equitable, and effective integration. Educational Management Administration & Leadership. https://doi.org/10.1177/17411432261418940

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. American Sociological Review, 48(2), 147-160. https://doi.org/10.2307/2095101

Donthu, N., Kumar, S., Mukherjee, D., Pandey, N., & Lim, W. M. (2021). How to conduct a bibliometric analysis: An overview and guidelines. Journal of Business Research, 133, 285-296. https://doi.org/10.1016/j.jbusres.2021.04.070

Doolin, B., Techatassanasoontorn, A. A., Waizenegger, L., & Wallace-Carter, E. (2025). Theorising robotic process automation as socio-technical change: A process study. Australasian Journal of Information Systems, 29, Article 5451. https://doi.org/10.3127/ajis.v29.5451

Dunleavy, P., Margetts, H., Bastow, S., & Tinkler, J. (2006). New public management is dead-Long live digital-era governance. Journal of Public Administration Research and Theory, 16(3), 467-494. https://doi.org/10.1093/jopart/mui057

Durao, D., & dos Reis, A. P. (2024). How does robotic process automation create value for firms? Information Systems and e-Business Management, 22(4), 721-740. https://doi.org/10.1007/s10257-024-00685-z

Fuszl, A., Zatorska, B., Van den Nest, M., Ebner, J., Presterl, E., & Diab-Elschahawi, M. (2021). The use of a UV-C disinfection robot in the routine cleaning process: A field study in an academic hospital. Antimicrobial Resistance & Infection Control, 10, Article 84. https://doi.org/10.1186/s13756-021-00945-4

Germundsson, N., Stranz, H., & Bergmark, A. (2024). Reducing administration? Examining the alignment of robotic process automation and social assistance in Swedish personal social services. Nordic Social Work Research. https://doi.org/10.1080/2156857X.2024.2440720

Gorman, A., & Brennan, A. (2026). Leading artificial intelligence amid uncertainty: Principals' sensemaking in Irish primary schools. Management in Education. https://doi.org/10.1177/08920206261468463

Grant, M. J., & Booth, A. (2009). A typology of reviews: An analysis of 14 review types and associated methodologies. Health Information and Libraries Journal, 26(2), 91-108. https://doi.org/10.1111/j.1471-1842.2009.00848.x

Guarascio, D., Piccirillo, A., & Reljic, J. (2025). Robots vs. workers: Evidence from a meta-analysis. Journal of Economic Surveys, 39(5), 2254-2271. https://doi.org/10.1111/joes.12699

Guarda, T., Sanchez, J., Augusto, M. F., & Torres, W. (2025). The impact of robotic process automation in higher education. In Advanced research in technologies, information, innovation and sustainability (CCIS Vol. 2346, pp. 19-30). Springer. https://doi.org/10.1007/978-3-031-83210-9_2

Gursoy, D., Chi, O. H., Lu, L., & Nunkoo, R. (2019). Consumers' acceptance of artificially intelligent (AI) device use in service delivery. International Journal of Information Management, 49, 157-169. https://doi.org/10.1016/j.ijinfomgt.2019.03.008

Herd, P., & Moynihan, D. P. (2018). Administrative burden: Policymaking by other means. Russell Sage Foundation.

Horpynich, H., Mistry, T. G., & Dogan, S. (2025). Service robots in hospitality: A cognitive appraisal perspective on job insecurity, turnover intentions, and generational differences. Journal of Hospitality and Tourism Technology, 16(1), 194-212. https://doi.org/10.1108/JHTT-03-2024-0189

Huynh, D.-M.-Q., & Lindgren, I. (2026). Beyond the hype of robotic process automation (RPA): On conditions needed to implement RPA in organizations. Information Systems and e-Business Management. https://doi.org/10.1007/s10257-026-00721-0

Jiang, X., & Abdullah, Z. (2026). AI-enabled governance in higher education: A systematic review of applications, outcomes, and emerging implications. Frontiers in Education, 11, Article 1856440. https://doi.org/10.3389/feduc.2026.1856440

Juell-Skielse, G., Guner, E. O., & Han, S. (2022). Adoption of robotic process automation in the public sector: A survey study in Sweden. In Electronic Government (EGOV 2022) (LNCS Vol. 13391, pp. 336-352). Springer. https://doi.org/10.1007/978-3-031-15086-9_22

Khantong, S., & Sriboonlue, P. (2026). Robotic process automation in business process management: A systematic literature review and an integrated framework. Technologies, 14(4), Article 225. https://doi.org/10.3390/technologies14040225

Kraus, P., Fissler, E., & Schlegel, D. (2024). A typology of challenges in the context of robotic process automation implementation projects. International Journal of Organizational Analysis, 32(11), 60-73. https://doi.org/10.1108/IJOA-11-2023-4100

Leonardi, P. M. (2011). When flexible routines meet flexible technologies: Affordance, constraint, and the imbrication of human and material agencies. MIS Quarterly, 35(1), 147-167. https://doi.org/10.2307/23043493

Lindgren, I., Johansson, B., Soderstrom, F., & Toll, D. (2022). Why is it difficult to implement robotic process automation? In Electronic Government (EGOV 2022) (LNCS Vol. 13391, pp. 353-368). Springer. https://doi.org/10.1007/978-3-031-15086-9_23

Lipsky, M. (2010). Street-level bureaucracy: Dilemmas of the individual in public services (30th anniversary expanded ed.). Russell Sage Foundation.

Lipsou, E., Keravnos, N., & Eteokleous, N. (2026). Artificial intelligence in school leadership: A structured literature review of organisational benefits and ethical challenges. Artificial Intelligence in Education, 2(3), 85-100. https://doi.org/10.1108/AIIE-08-2025-0240

Macalupu, V., Miller, E., Martin, L., & Caldwell, G. (2025). Human-robot interactions and experiences of staff and service robots in aged care. Scientific Reports, 15, Article 2495. https://doi.org/10.1038/s41598-025-86255-w

Marchena Sekli, G. F., & Godo, A. (2025). Transforming teaching and learning with robotic process automation: A systematic review of pedagogical applications. Journal of Information Technology Education: Research, 24, Article 39. https://doi.org/10.28945/5655

Marcu, G., Lin, I., Williams, B., Robert, L. P., Jr., & Schaub, F. (2023). “Would I feel more secure with a robot?”: Understanding perceptions of security robots in public spaces. Proceedings of the ACM on Human-Computer Interaction, 7(CSCW2), 1-34. https://doi.org/10.1145/3610171

Marquez, L., Henriquez, V., Chevreux, H., Scheihing, E., & Guerra, J. (2024). Adoption of learning analytics in higher education institutions: A systematic literature review. British Journal of Educational Technology, 55(2), 439-459. https://doi.org/10.1111/bjet.13385

Meyer, J. W., & Rowan, B. (1977). Institutionalized organizations: Formal structure as myth and ceremony. American Journal of Sociology, 83(2), 340-363. https://doi.org/10.1086/226550

Misaros, M., Stan, O. P., Enyedi, S., Stan, A., Donca, I., & Miclea, L. C. (2024). A method for assessing the reliability of the Pepper robot in handling office documents: A case study. Biomimetics, 9(9), Article 558. https://doi.org/10.3390/biomimetics9090558

O'Dea, M., Zhou, X., Teng, D., Mundy, D., & Ishaya, T. (2024). Are technology acceptance models still fit for purpose? [Editorial]. Journal of University Teaching and Learning Practice, 21(8). https://doi.org/10.53761/1bdbms32

Ohneberg, C., Stobich, N., & Eberl, I. (2025). Possible applications and conditions for the development and integration of robotic assistance systems in the acute inpatient setting from the perspective of nursing professionals: A qualitative study. International Journal of Social Robotics, 17(9), 1821-1835. https://doi.org/10.1007/s12369-025-01289-4

Oladokun, B. D., Maxwell, C. E., Umar, L., Alex-Nmecha, J. C., & Adigun, G. O. (2026). The rise of robots in libraries: Essential skills and knowledge for librarians in Industry 5.0. Information Services and Use, 46(3), 47-57. https://doi.org/10.1177/18758789261443981

Oliveira, T., & Martins, M. F. (2011). Literature review of information technology adoption models at firm level. The Electronic Journal of Information Systems Evaluation, 14(1), 110-121.

Oruma, S. O., Ayele, Y. Z., Sechi, F., & Rodsethol, H. (2023). Security aspects of social robots in public spaces: A systematic mapping study. Sensors, 23(19), Article 8056. https://doi.org/10.3390/s23198056

Oruma, S. O., Colomo-Palacios, R., & Gkioulos, V. (2026). From framework to reliable practice: End-user perspectives on social robots in public spaces. Systems, 14(2), Article 137. https://doi.org/10.3390/systems14020137

Ozturkcan, S. (2025). Responsible robot integration for inclusive and sustainable work. Discover Sustainability, 6, Article 912. https://doi.org/10.1007/s43621-025-01928-w

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. BMJ, 372, n71. https://doi.org/10.1136/bmj.n71

Paluch, S., Tuzovic, S., Holz, H. F., Kies, A., & Jorling, M. (2021). “My colleague is a robot” - Exploring frontline employees' willingness to work with collaborative service robots. Journal of Service Management, 33(2), 363-388. https://doi.org/10.1108/JOSM-11-2020-0406

Polat, M., Karatas, I. H., & Varol, N. (2025). Ethical artificial intelligence (AI) in educational leadership: Literature review and bibliometric analysis. Leadership and Policy in Schools, 24(1), 46-76. https://doi.org/10.1080/15700763.2024.2412204

Rogers, E. M. (2003). Diffusion of innovations (5th ed.). Free Press.

Rughinis, C., Vulpe, S.-N., Turcanu, D., & Rughinis, R. (2025). AI at the knowledge gates: Institutional policies and hybrid configurations in universities and publishers. Frontiers in Computer Science, 7, Article 1608276. https://doi.org/10.3389/fcomp.2025.1608276

Sarker, S., Chatterjee, S., Xiao, X., & Elbanna, A. (2019). The sociotechnical axis of cohesion for the IS discipline: Its historical legacy and its continued relevance. MIS Quarterly, 43(3), 695-719. https://doi.org/10.25300/MISQ/2019/13747

Shahzad, K., Khan, S. A., & Iqbal, A. (2025). Factors influencing the adoption of robotic technologies in academic libraries: A systematic literature review (SLR). Journal of Librarianship and Information Science, 57(3), 687-704. https://doi.org/10.1177/09610006241231012

Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. Journal of Business Research, 104, 333-339. https://doi.org/10.1016/j.jbusres.2019.07.039

Swist, T., & Gulson, K. N. (2022). School choice algorithms: Data infrastructures, automation, and inequality. Postdigital Science and Education, 5(1), 152-170. https://doi.org/10.1007/s42438-022-00334-z

Syed, R., Suriadi, S., Adams, M., Bandara, W., Leemans, S. J. J., Ouyang, C., ... Reijers, H. A. (2020). Robotic process automation: Contemporary themes and challenges. Computers in Industry, 115, Article 103162. https://doi.org/10.1016/j.compind.2019.103162

Tella, A., & Ajani, Y. A. (2022). Robots and public libraries. Library Hi Tech News, 39(7), 15-18. https://doi.org/10.1108/LHTN-05-2022-0072

Toma, M.-V. (2023). Robotic process automation in higher education: A case study on university daily sheets. Journal of Applied Computer Science & Mathematics, 17(2), 21-26. https://doi.org/10.4316/JACSM.202302003

Tornatzky, L. G., Fleischer, M., & Chakrabarti, A. K. (1990). The processes of technological innovation. Lexington Books.

Torraco, R. J. (2005). Writing integrative literature reviews: Guidelines and examples. Human Resource Development Review, 4(3), 356-367. https://doi.org/10.1177/1534484305278283

UiPath. (n.d.). Use cases for RPA in education - University of Auckland [Vendor case study]. https://www.uipath.com/resources/automation-case-studies/university-of-auckland

Uta, B.-I., Dascalu, M.-I., Neagu, A.-M., Guica, R. I., & Teodorescu, I.-E. (2026). Supporting educational administration via emergent technologies: A case study for a faculty of engineering in foreign languages. Education Sciences, 16(1), Article 29. https://doi.org/10.3390/educsci16010029

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. MIS Quarterly, 27(3), 425-478. https://doi.org/10.2307/30036540

Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. MIS Quarterly, 36(1), 157-178. https://doi.org/10.2307/41410412

Venkatesh, V., Thong, J. Y. L., & Xu, X. (2016). Unified theory of acceptance and use of technology: A synthesis and the road ahead. Journal of the Association for Information Systems, 17(5), 328-376.

Volkoff, O., & Strong, D. M. (2013). Critical realism and affordances: Theorizing IT-associated organizational change processes. MIS Quarterly, 37(3), 819-834. https://doi.org/10.25300/MISQ/2013/37.3.07

Wang, Y. (2024). Algorithmic decisions in education governance: Implications and challenges. Discover Education, 3(1), 1-10. https://doi.org/10.1007/s44217-024-00337-x

Weinberg, D., Dwyer, H., Fox, S. E., & Martelaro, N. (2023). Sharing the sidewalk: Observing delivery robot interactions with pedestrians during a pilot in Pittsburgh, PA. Multimodal Technologies and Interaction, 7(5), Article 53. https://doi.org/10.3390/mti7050053

Willcocks, L., Lacity, M., & Craig, A. (2017). Robotic process automation: Strategic transformation lever for global business services? Journal of Information Technology Teaching Cases, 7(1), 17-28. https://doi.org/10.1057/s41266-016-0016-9

Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2018). Brave new world: Service robots in the frontline. Journal of Service Management, 29(5), 907-931. https://doi.org/10.1108/JOSM-04-2018-0119

Ye, X., & Robert, L. P. (2024). A human-security robot interaction literature review. ACM Transactions on Human-Robot Interaction, 14(2), 1-36. https://doi.org/10.1145/3700888

Young, M. M., Bullock, J. B., & Lecy, J. D. (2019). Artificial discretion as a tool of governance: A framework for understanding the impact of artificial intelligence on public administration. Perspectives on Public Management and Governance, 2(4), 301-313. https://doi.org/10.1093/ppmgov/gvz014

Zeng, M., Cheah, K. S. L., & Abdullah, Z. (2025). The influence of school principals' digital leadership on teachers' competency in integrating artificial intelligence: A systematic thematic review. Frontiers in Education, 10, Article 1655967. https://doi.org/10.3389/feduc.2025.1655967