Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Role of Artificial Intelligence and the Internet of Things in Education and Cultural Studies

Arun Kumar Singh

arun.singh@pnguot.ac.pg

PNG University of Technology

Research Article

Keywords: Artificial Intelligence in Education, Internet of Things (IoT), Learning Analytics, Algorithmic Fairness, Cultural Heritage, Digital Humanities, Data Governance

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11605123/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Role of Artificial Intelligence and the Internet of Things in Education and Cultural Studies

Arun Kumar Singh

School of MCS, PNG University of Technology, Lae, Papua New Guinea,

ORCID Id: 0000-0002-9974-4696, E-mail: arun.singh@pnguot.ac.pg, arunsinghiiita@gmail.com, Phone No: +67581506607, + 917895276828

Abstract: AI and IoT add most value in education and cultural studies if they expand who is reached and protected, and our simulations suggest that they may narrow that reach almost as easily. This paper (i) synthesizes the literature on AI (Artificial Intelligence) and the Internet of Things (IoT) across teaching and learning, assessment, campus operations, digital humanities, and heritage preservation, and (ii) evaluates two representative applications with reproducible Monte Carlo models. The first is an early-warning system that directs limited tutoring capacity to at-risk students, with an explicit equity metric for a minority group whose data are less informative. The second is IoT microclimate monitoring in a heritage gallery with three alarm strategies. In a simulated cohort of 2,000 students (15% at risk, 300 tutoring slots), random allocation reached 15.2% of at-risk students, LMS analytics reached 37.4% and combining LMS and IoT data reached 50.2%. However, the minority group was reached 21.7 percentage points less often than the majority group. Allocating capacity by group quotas cut that gap to 15.7 points at the same overall reach (50.0%) and only equal data quality removed it. In the gallery simulation, a fixed humidity threshold detected 47.0% of excursions and raised 12.4 false alarms per 30 days, whereas a three-sensor consensus rule detected 98.0% with 0.7 false alarms. We conclude that data quality and fairness, not model accuracy alone, determine educational benefit, and that sensor redundancy matters more than a single smarter detector. We propose a four-principle governance framework covering equity, privacy, cultural sovereignty, and human oversight. All results are simulation-based and the full Python code is provided.

Keywords: Artificial Intelligence in Education, Internet of Things (IoT), Learning Analytics, Algorithmic Fairness, Cultural Heritage, Digital Humanities, Data Governance

1. Introduction

Artificial intelligence (AI) and the Internet of Things (IoT) are now a days embedded in classrooms, libraries, museums, and archives, and the question for scholars has changed from whether they matter to for whom they matter and at what cost.

Background and motivation

Education has been a proving ground for AI for decades, and the field has exploded. A systematic review of AI applications in higher education found that research was dominated by computer-science perspectives, that educators were rarely consulted, and that critical

reflection on risks and implications for pedagogical theory were lacking (Zawacki-Richter et al., 2019). Subsequent reviews cover the breadth of applications, from intelligent tutoring and automatic assessment to adaptive content and administrative support (Chen et al., 2020; Hwang et al., 2020; Zhang & Aslan, 2021), while critical scholars highlight that the promises of AI in education repeat older patterns of over-claiming (Williamson & Eynon, 2020; Selwyn, 2022).

Learning analytics is the most developed data-driven strand. Early work showed that activity data from learning management systems (LMS) can be used for early-warning systems for educators (Macfadyen & Dawson, 2010), and machine-learning frameworks have been proposed to identify students at risk of adverse academic outcomes (Lakkaraju et al., 2015). Overviews of educational data mining and learning analytics track how methods, aims, and concerns have evolved (Ferguson, 2012; Romero & Ventura, 2020). Since late 2022, generative AI has been a second strand, with large language models (LLMs) being able to draft, explain, and tutor, with documented opportunities and challenges for teaching, learning, and academic integrity (Kasneci et al., 2023; Cotton et al., 2024), and controlled evidence outside education indicates significant productivity effects on writing tasks (Noy & Zhang, 2023).

The IoT adds a physical dimension. The IoT connects sensors, tags, and actuators to networks so that objects can be monitored and controlled (Atzori et al., 2010; Gubbi et al., 2013; Al-Fuqaha et al., 2015). A systematic review of IoT in education finds benefits for engagement, resource management, and campus operations, alongside challenges of costs, interoperability, privacy, and security (Kassab et al., 2020). Security and privacy concerns in distributed IoT deployments are well-documented (Roman et al., 2013; Sicari et al., 2015).

Cultural studies and heritage use the same technologies in different ways. Machine learning supports the analysis, classification, and restoration of cultural heritage materials (Fiorucci et al., 2020), augmented, virtual, and mixed reality extends access to sites and collections (Bekele et al., 2018), and environmental sensing supports preventive conservation at a time when climate change is threatening heritage assets (Sesana et al., 2021). At the same time, cultural scholars have raised a distinct concern: the models on which tools are based encode particular cultural values. A recent evaluation found that the values expressed by widely used LLMs are similar to those of English-speaking and Protestant European countries (Tao et al., 2024), and detectors of AI-written text have been shown to misclassify writing by non-native English speakers (Liang et al., 2023).

Problem statement and research gap

Three gaps motivate this paper. First, AI and IoT are usually reviewed separately, and education and cultural studies are rarely together, despite both relying on collecting, interpreting, and preserving human and material records. Second, claims of benefit rarely

report who is left out. Algorithmic bias in education and it is well-documented (Baker & Hawn, 2022), and a widely-cited healthcare example shows that a seemingly neutral prediction target can be biased (Obermeyer et al., 2019), while planning for AI-supported student support rarely quantifies the equity cost of unequal data quality. Third part, in the heritage setting, the practical question is not only whether sensors and algorithms work but which alarm design is reliable when sensors flop or problem.

Research questions

- RQ1. AI and IoT currently contribute across education and cultural studies, and what benefits and risks does the evidence report, where do?
- RQ2. How much of limited intervention capacity can AI-based early warning direct to at-risk students, and how does data quality for a minority group change who is reached?
- RQ3. Alternative alarm strategies for IoT microclimate monitoring compare on detection, delay, and false alarms, how do?
- RQ4. Principles follow from the answers to RQ1–RQ3, Which governance?

Contributions and Organization

This paper describes a structured synthesis across six application domains, two open Python programs that make assumptions explicit, an equity-aware evaluation of early-warning allocation, a comparison of three IoT alarm strategies under sensor glitches, and a four-principle governance framework. Section 2 describes materials and methods. Section 3 presents and discusses results and limitations. Section 4 concludes with recommendations.

2. Materials and Methods

The study combines a structured literature synthesis (RQ1) with two Monte Carlo evaluations (RQ2 and RQ3) and uses all three to derive governance principles (RQ4).

Study design

Figure 1 shows the pipeline. The synthesis organises evidence into six application domains. Two domains were then selected for quantitative evaluation because they pose a decision that can be modelled: how to allocate scarce student-support capacity (student success and analytics) and how to design alarms for sensor networks (heritage preservation). The simulations make assumptions explicit and show how outcomes change when those assumptions change.

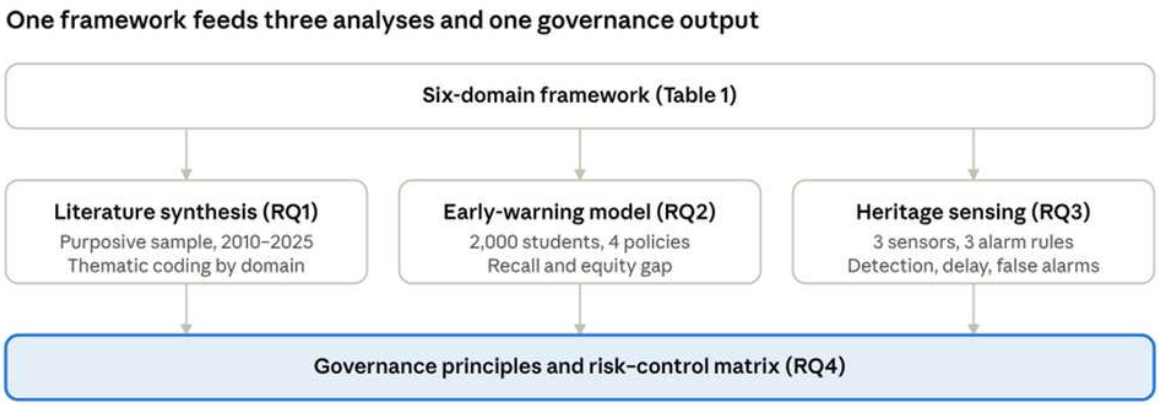


Figure 1. Study pipeline: one domain framework feeds a synthesis and two simulations, which together yield governance principles.

Literature synthesis

The synthesis is narrative and purposive, not an exhaustive systematic review. Sources are peer-reviewed articles, reviews, and policy-relevant commentaries published mostly between 2010 and 2025 that address AI or IoT in education, digital humanities, or heritage. They were selected for relevance, citation visibility, and coverage of each domain. We applied PRISMA 2020 transparency principles where applicable (Page et al., 2021) but did not run a formal multi-database search or produce a flow diagram. Each source was coded to a domain and to a benefit or risk theme using reflexive thematic analysis (Braun & Clarke, 2006). Table 1 gives the coding frame.

Table 1. Coding frame: six application domains.

Domain	Example tasks	AI role	IoT role
1. Teaching and learning	Tutoring, feedback, adaptive content	LLM tutors, recommenders	Classroom and lab sensing
2. Assessment and integrity	Scoring support, integrity checks	Automated scoring, AI-text detection	Exam-room monitoring devices
3. Student success and analytics	Early warning, retention support	Predictive models	Attendance, access, and usage sensors
4. Campus and library operations	Space, energy, and access management	Demand forecasting	Occupancy and environment sensors
5. Digital humanities and cultural analysis	Text and image analysis, translation	Natural language processing, computer vision	Digitisation and tagging devices

Domain	Example tasks	AI role	IoT role
6. Heritage preservation and museums	Conservation, visitor experience	Anomaly detection, immersive content	Microclimate and vibration sensors

Experiment A: early-warning allocation and equity (RQ2)

Here a cohort of N = 2,000 students includes a minority group (20%) whose data are less informative for the model, aimed at example because of language, connectivity, or sensor coverage. Fifteen percent of students are at risk. Staff can offer support to K = 300 students. A model assigns each student a risk score drawn from a binormal model: scores for non-risk students follow N(0, 1), and scores for at-risk students are shifted by a group-specific amount chosen to reach a target AUC (Equation 1).

$$s_i = \varepsilon_i + y_i\,\mu_{g(i)},\qquad \varepsilon_i \sim \mathcal{N}(0,1),\qquad \mu_g = \sqrt{2}\,\Phi^{-1}(\mathrm{AUC}_g)$$

Equation 1. Binormal score model; $y_i = 1$ for an at-risk student and g(i) is the student's group.

Four policies are compared. **P0** allocates the 300 slots at random. **P1** ranks students by an LMS-only model (AUC 0.75 for the majority, 0.68 for the minority) and supports the top 300. **P2** adds IoT-derived features such as campus-access and device-usage signals (AUC 0.85 and 0.74) and again takes the top 300. **P3** uses the P2 scores but splits capacity between groups in proportion to group size and ranks within each group. The lower minority AUC values reflect the documented tendency of models to perform worse for groups that are under-represented in training data (Baker & Hawn, 2022; Mehrabi et al., 2021). The equity metric is the gap in recall between groups, which corresponds to the equal-opportunity criterion (Hardt et al., 2016). Outcomes are in Equation 2.

$$\mathrm{Recall}_g = \frac{\#\{\text{at-risk flagged in } g\}}{\#\{\text{at-risk in } g\}},\qquad \mathrm{Gap} = \mathrm{Recall}_{maj} - \mathrm{Recall}_{min},\qquad \Delta R = u \cdot \mathrm{TP}$$

Equation 2. Group recall, equal-opportunity gap, and expected additional students retained, where TP is the number of at-risk students flagged and u = 0.15 is an assumed intervention effect.

Table 2. Parameters of Experiment A.

Parameter	Baseline value	Range explored	Role
Cohort size	2,000	fixed	Students per replication
Minority share	0.20	fixed	Group with less informative data
At-risk prevalence	0.15	fixed	Same in both groups
Support capacity K	300	fixed	Slots equal to 15% of the cohort
AUC, LMS analytics (maj. / min.)	0.75 / 0.68	fixed	Policy P1
AUC, LMS + IoT (maj. / min.)	0.85 / 0.74	minority 0.65–0.85	Policies P2 and P3
Intervention effect u	0.15	fixed	Share of reached at-risk students retained
Replications	1,000	300 in sensitivity	Monte Carlo sample size

Experiment B: IoT microclimate monitoring in a heritage gallery (RQ3)

There are three collocated humidity sensors sample every 10 minutes for 36 days: 6 days for warming up and training, and 30 for evaluation. The accurate room humidity exhibits a diurnal cycle plus an autoregressive disturbance. Four excursions per month, for instance, caused by a ventilation failure, raise the humidity by 4-12 percentage points (with a ramp, a hold and a decay) (Figure 3 reports detection by excursion size). Each sensor adds measurement noise and occasional single-sample spikes which mimic hardware glitches. Preventive conservation is interested in detecting true excursions early, and the threat of climate-driven deterioration makes the problem compelling (Sesana et al., 2021). Three alarm strategies are compared: D1 alarms when any sensor exceeds a fixed 60% relative-humidity threshold. D2 applies an exponentially weighted moving-average (EWMA) control chart (Lucas & Saccucci, 1990) to one sensor's residuals from the same-time-of-day median of the previous three days (Equation 3). D3 is alarmed only when at least two of the three sensors' EWMA charts alarm simultaneously. Anomaly detection in general is reviewed by Chandola et al. (2009).

$$r_t = x_t - \operatorname{median}(x_{t-D},\,x_{t-2D},\,x_{t-3D}),\qquad z_t = \lambda\, r_t + (1-\lambda)\, z_{t-1},\qquad \text{alarm if } |z_t| > k\,\sigma_z$$

Equation 3. EWMA chart on seasonal residuals; D is one day of samples, λ = 0.2, k = 4, and σz is estimated robustly from the training days.

For each detector we record the share of excursions detected within the excursion window, the mean delay from onset to first alarm among detected excursions, and the number of false-alarm episodes per 30 days (alarms outside excursion windows, merged if less than an hour apart).

Table 3. Parameters of Experiment B.

Parameter	Baseline value	Range explored
Sampling interval, sensors	10 min, 3 sensors	fixed
Period	6 warm-up days + 30 evaluation days	fixed
Baseline humidity, diurnal amplitude	50% RH, 3% RH	fixed
Disturbance, sensor noise	AR(1) 0.9 with innovation SD 0.35; SD 0.5	fixed
Glitch probability per sample per sensor	0.002 (spike of +10% RH)	0.0005–0.008
Excursions per 30 days, size	4, uniform 4–12% RH	fixed
Excursion shape	2 h ramp, 1 h hold, 2 h decay	fixed
D1 threshold	60% RH	fixed
D2 and D3 chart	$\lambda = 0.2$, 4σ limit; D3 needs 2 of 3 sensors	fixed
Replications	300	100 in glitch sensitivity

Implementation

Both experiments are implemented in Python 3 with NumPy, pandas, and SciPy. The listing below is complete, and running it prints every table reported in Section 3. Seeds are fixed, so results are exactly reproducible.

```
“””
Monte Carlo evaluation for the paper “Role of AI and IoT in Education and Cultural Studies”.
Experiment A: AI/IoT early-warning analytics for at-risk students, with a group-equity metric.
Experiment B: IoT microclimate monitoring of a heritage gallery with three alarm strategies.
All data are simulated; every parameter is a stated assumption (see paper, Tables 2 and 3).
“””
import numpy as np
import pandas as pd
from scipy.stats import norm
from scipy.signal import lfilter

SEED = 42

# ---------------- Experiment A: early-warning analytics ----------------
```

```
N_STUD, MINORITY_SHARE, PREV = 2000, 0.20, 0.15
CAPACITY, UPLIFT, N_REP_A = 300, 0.15, 1000
POLICIES = {                    # (policy kind, AUC majority, AUC minority)
    “P0 random allocation”:       (“random”, 0.50, 0.50),
    “P1 LMS analytics”:           (“top”,    0.75, 0.68),
    “P2 LMS + IoT”:               (“top”,    0.85, 0.74),
    “P3 LMS + IoT, group quotas”: (“quota”,  0.85, 0.74),
}


def mu(auc):
    “””Mean shift of a unit-variance binormal score model that yields a given AUC.”””
    return np.sqrt(2.0) * norm.ppf(auc)


def pick(rng, s, g, kind, k):
    n = s.size
    if kind == “random”:
        return rng.choice(n, k, replace=False)
    if kind == “quota”:           # capacity split in proportion to group size
        out = []
        for grp in (True, False):
            members = np.flatnonzero(g == grp)
            kk = int(round(k * members.size / n))
            out.extend(members[np.argsort(-s[members])[:kk]])
        return np.array(out)
    return np.argsort(-s)[:k]


def run_a(kind, auc_maj, auc_min, k=CAPACITY, n_rep=N_REP_A, seed=SEED):
    rng = np.random.default_rng(seed)
    rows = []
    for _ in range(n_rep):
        g = rng.random(N_STUD) < MINORITY_SHARE          # True = minority group
        y = rng.random(N_STUD) < PREV                    # True = at risk
        s = rng.normal(size=N_STUD) + y * np.where(g, mu(auc_min), mu(auc_maj))
        flag = np.zeros(N_STUD, bool)
        flag[pick(rng, s, g, kind, k)] = True
        tp = (flag & y).sum()
        rec_maj = (flag & y & ~g).sum() / max((y & ~g).sum(), 1)
        rec_min = (flag & y & g).sum() / max((y & g).sum(), 1)
        rows.append((tp / max(y.sum(), 1), rec_maj, rec_min, rec_maj – rec_min,
```

```
                tp / flag.sum(), tp * UPLIFT))
    return pd.DataFrame(rows, columns=[“recall”, “recall_maj”, “recall_min”,
                                       “gap”, “precision”, “extra_retained”])


def table_a():
    out = {}
    for name, (kind, am, an) in POLICIES.items():
        d = run_a(kind, am, an)
        out[name] = dict(recall=d.recall.mean(), lo=d.recall.quantile(.025),
                         hi=d.recall.quantile(.975), recall_maj=d.recall_maj.mean(),
                         recall_min=d.recall_min.mean(), gap=d.gap.mean(),
                         precision=d.precision.mean(), extra=d.extra_retained.mean())
    return pd.DataFrame(out).T.round(3)


def sens_minority_auc(values=(0.65, 0.70, 0.74, 0.80, 0.85), n_rep=300):
    rows = []
    for an in values:
        p2 = run_a(“top”, 0.85, an, n_rep=n_rep)
        p3 = run_a(“quota”, 0.85, an, n_rep=n_rep)
        rows.append((an, p2.gap.mean(), p3.gap.mean(), p2.recall.mean(), p3.recall.mean()))
    return pd.DataFrame(rows, columns=[“auc_min”, “gap_P2”, “gap_P3”,
                                       “recall_P2”, “recall_P3”]).round(3)


# ---------------- Experiment B: heritage microclimate sensing ----------------
D = 144                                 # samples per day (10-minute sampling)
DAYS, WARM_DAYS = 36, 6                 # 6 warm-up/training days + 30 evaluation days
T, EVAL0 = DAYS * D, WARM_DAYS * D
N_SENS, EVENTS = 3, 4
BASE_RH, DIURNAL, AR, AR_SD, SENS_SD = 50.0, 3.0, 0.9, 0.35, 0.5
GLITCH_P, GLITCH_MAG = 0.002, 10.0      # isolated single-sample sensor spikes
MAG_RANGE = (4.0, 12.0)                 # excursion size in % relative humidity
SHAPE = np.r_[np.linspace(0, 1, 13)[1:], np.ones(6), np.linspace(1, 0, 13)[1:]]
LEN, EXT = len(SHAPE), len(SHAPE) + 6
FIXED_THR, LAMBDA, K_SIGMA = 60.0, 0.2, 4.0
SAMPLE_MIN, N_REP_B = 10, 300


def gallery(rng):
    t = np.arange(T)
```

```
    e = np.zeros(T)
    eps = rng.normal(0, AR_SD, T)
    for I in range(1, T):
        e[i] = AR * e[I – 1] + eps[i]
    room = BASE_RH + DIURNAL * np.sin(2 * np.pi * t / D) + e
    starts = []
    while len(starts) < EVENTS:
        s = int(rng.integers(EVAL0 + 12, T – LEN – 12))
        if all(abs(s – x) > 2 * D for x in starts):
            starts.append(s)
    starts = sorted(starts)
    mags = rng.uniform(*MAG_RANGE, EVENTS)
    for s, m in zip(starts, mags):
        room[s:s + LEN] += m * SHAPE
    x = room + rng.normal(0, SENS_SD, (N_SENS, T)) \
        + GLITCH_MAG * (rng.random((N_SENS, T)) < GLITCH_P)
    return x, starts, mags


def ewma_alarm(x):
    “””EWMA chart on residuals from the same-time-of-day median of the previous 3 days.”””
    base = np.full(T, np.nan)
    base[3 * D:] = np.median(np.vstack([x[2 * D:-D], x[D:-2 * D], x[:-3 * D]]), axis=0)
    r = np.nan_to_num(x – base)
    z = lfilter([LAMBDA], [1, -(1 – LAMBDA)], r)
    tr = z[3 * D:6 * D]
    sig = 1.4826 * np.median(np.abs(tr – np.median(tr)))
    return np.abs(z) > K_SIGMA * sig


def evaluate(alarm, starts):
    detected, delay = [], []
    for s in starts:
        w = alarm[s:s + EXT]
        detected.append(bool(w.any()))
        delay.append(np.argmax(w) * SAMPLE_MIN if w.any() else np.nan)
    out = alarm.copy()
    out[:EVAL0] = False
    for s in starts:
        out[s:s + EXT] = False
    idx = np.flatnonzero(out)
    false_alarms = 0 if idx.size == 0 else 1 + int((np.diff(idx) > 6).sum())
```

```
    return detected, delay, false_alarms


def run_b(n_rep=N_REP_B, seed=SEED):
    rng = np.random.default_rng(seed)
    ev, fa = [], []
    for _ in range(n_rep):
        x, starts, mags = gallery(rng)
        per = np.array([ewma_alarm(x[i]) for I in range(N_SENS)])
        alarms = {“D1 fixed threshold”: (x > FIXED_THR).any(axis=0),
                  “D2 single-sensor EWMA”: per[0],
                  “D3 three-sensor consensus”: per.sum(axis=0) >= 2}
        res = {k: evaluate(a, starts) for k, a in alarms.items()}
        for j, m in enumerate(mags):
            for k, (det, dly, _) in res.items():
                ev.append((k, m, det[j], dly[j]))
        for k, (_, _, f) in res.items():
            fa.append((k, f))
    return (pd.DataFrame(ev, columns=[“detector”, “mag”, “detected”, “delay_min”]),
            pd.DataFrame(fa, columns=[“detector”, “false_alarms”]))


def table_b(ev, fa):
    g = ev.groupby(“detector”)
    out = pd.DataFrame({“detection_rate_pct”: 100 * g.detected.mean(),
                        “mean_delay_min”: g.delay_min.mean(),
                        “false_alarms_per_30d”: fa.groupby(“detector”).false_alarms.mean()})
    return out.round(2)


def detection_by_magnitude(ev):
    ev = ev.assign(bin=np.floor(ev.mag).astype(int).clip(4, 11))
    return (100 * ev.pivot_table(index=”bin”, columns=”detector”,
                                 values=”detected”, aggfunc=”mean”)).round(1)


def sens_glitch(values=(0.0005, 0.002, 0.008), n_rep=100):
    “””False alarms per 30 days for each detector as the sensor glitch rate varies.”””
    global GLITCH_P
    base, rows = GLITCH_P, []
    try:
        for p in values:
```

```
            GLITCH_P = p
            ev, fa = run_b(n_rep=n_rep)
            rows.append((p, *table_b(ev, fa).false_alarms_per_30d.values))
    finally:
        GLITCH_P = base
    return pd.DataFrame(rows, columns=[“glitch_rate”, “D1_fixed”, “D2_single”,
“D3_consensus”])


if  == ″__ma__″“
   d.se”_option("display.wid“h", 200)
    ”d.set_option("display.max“columns", 20)
    p”int(table_a())
    print(sens_minority_auc())
    ev, fa = run_b()
    print(table_b(ev, fa))
    print(detection_by_magnitude(ev))
    print(sens_glitch())
```

Replication, uncertainty, and ethics

Baseline results are based on one thousand replications for Experiment A and three hundred for Experiment B, and sensitivity runs are based on three hundred and one hundred replications. Intervals for recall are the 2.5th and 97.5th percentiles over replications. Since there were no human participants or personal data involved, ethical review was not required. In practice, student analytics and sensor data pose consent, privacy, and security risks to participants (Slade & Prinsloo, 2013; Roman et al., 2013) and are deeply intertwined with community data-governance principles (Carroll et al., 2020). In order to conform to the Findable, Accessible, Interoperable, and Reusable (FAIR) data principles (Wilkinson et al., 2016), the code is self-contained and the parameters listed in Tables 2 and 3.

3. Results and Discussion

AI and IoT yield clear benefits in both fields, but in our models the benefits are unevenly distributed: richer data yielded benefits to the majority group about four times higher than to the minority group, and single sensors failed in ways easily compensated for by redundancy.

Evidence across the six domains (RQ1)

Table 4 clarify the evidence matrix by application domain. In this paper purposive sample, evidence of student-success analytics and of AI-assisted teaching is most

developed, evidence of IoT in education is predominantly descriptive, and this evidence of cultural-studies applications is scant and more technical than for education. Throughout domains the recurring risks are uneven data quality, inadequate privacy protection, and cultural value bias.

Table 4. Evidence matrix by application domain.

Domain	Strongest evidence	Main benefit	Main risk
1. Teaching and learning	Reviews of AI in education (Zawacki-Richter et al., 2019; Chen et al., 2020); generative AI (Kasneci et al., 2023; Noy & Zhang, 2023)	Personalised support; faster drafting and feedback	Over-reliance; fabricated content (Ji et al., 2023); cultural and linguistic bias (Tao et al., 2024; Bender et al., 2021)
2. Assessment and integrity	Academic-integrity analyses (Cotton et al., 2024)	Scalable feedback and scoring support	Unreliable AI-text detection; bias against non-native writers (Liang et al., 2023)
3. Student success and analytics	Early-warning systems (Macfadyen & Dawson, 2010; Lakkaraju et al., 2015); surveys (Romero & Ventura, 2020)	Earlier identification of at-risk students	Algorithmic bias (Baker & Hawn, 2022); consent and surveillance (Slade & Prinsloo, 2013)
4. Campus and library operations	Reviews of IoT in education (Kassab et al., 2020); IoT foundations (Gubbi et al., 2013)	Resource, space, and energy management	Cost, interoperability, security (Roman et al., 2013; Sicari et al., 2015)
5. Digital humanities and cultural analysis	Machine learning for heritage materials (Fiorucci et al., 2020)	Analysis of texts and images at scale	Encoded cultural values (Tao et al., 2024); fabricated attributions (Ji et al., 2023)
6. Heritage preservation and museums	Immersive technologies (Bekele et al., 2018); climate risk (Sesana et al., 2021)	Preventive conservation; wider access	False alarms; data sovereignty (Carroll et al., 2020)

In the teaching and learning, generative artificial intelligence (AI) provides the means to produce, explain, and critique at scale, raising issues about correctness and trust (Kasneci et al., 2023; Cotton et al., 2024). Typical works highlight the danger of technological solutionism (Selwyn, 2022; Williamson & Eynon, 2020), and surveys show that educational practice has lagged behind innovation (Zawacki-Richter et al., 2019). With in assessment,

the key finding on detectors of AI-written texts is that they are often biased against non-native English writers (Liang et al., 2023), making their use in classrooms an issue of fairness.

Aimed at student analytics, early warning systems are straight forward to build (Macfadyen & Dawson, 2010; Lakkaraju et al., 2015), but questions of bias, student consent, and whether or not institutions should act on these insights dominate the literature (Slade & Prinsloo, 2013; Baker & Hawn, 2022). An illustrative example from healthcare, where similar predictive models are used to allocate care, shows that a target value that seems reasonable can encode a bias that substantially hurt one group more than others (Obermeyer et al., 2019). Concerning the Internet of Things (IoT) on campus, surveys identify benefits in terms of engagement and management, but learning effectiveness is mixed (Kassab et al., 2020), and devices are often insecure (Roman et al., 2013; Sicari et al., 2015). Last but not least, for work in the humanities and cultural heritage, machine learning and immersive graphics enable new understandings, findings, and displays (Fiorucci et al., 2020; Bekele et al., 2018), while ethical considerations around community ownership of data are just being worked into practice (Carroll et al., 2020).

Early-warning allocation and equity (RQ2)

Adding IoT data lifted overall reach of at-risk students from 37.4% to 50.2%, but 15.1 of those percentage points went to the majority group and only 3.9 to the minority, so the equity gap doubled from 10.6 to 21.7 points (Table 5).

Table 5. Experiment A baseline: reach of at-risk students by policy (2,000 students, 15% at risk, 300 slots; 1,000 replications).

Policy	Overall recall (95% interval)	Majority recall	Minority recall	Gap (points)	Expected extra students retained
P0 random allocation	15.2% (11.6–19.0)	15.2%	15.2%	0.0	6.8
P1 LMS analytics	37.4% (32.7–41.9)	39.5%	28.9%	10.6	16.8
P2 LMS + IoT	50.2% (45.6–54.7)	54.6%	32.8%	21.7	22.6
P3 LMS + IoT, group quotas	50.0% (45.3–54.5)	53.2%	37.5%	15.7	22.5

Two features of Table 5 are worth noting. Firstly, the random allocation reaches 15.2% of at-risk students because the capacity is equivalent to 15% of the cohort, hence any model should be assessed against this level of recall rather than zero. Since the number of slots was the same as the expected number of at-risk students, precision equals recall in this design. Secondly, P2 achieves the greatest expected benefit of around 22.6 additional

students retained per cohort at the assumed 15% intervention effect compared to 16.8 for LMS-only analytics. Group quotas (P3) offer almost the same overall benefit (22.5) but raise minority recall from 32.8% to 37.5% and narrow the gap to 15.7 points at the expense of a small drop in majority recall, from 54.6% to 53.2%.

Figure 2 shows that the gap is primarily driven by data quality for the minority group rather than by the allocation rule itself.

ai_iot_edu_culture_sim.py · Experiment A · 300 replications per point

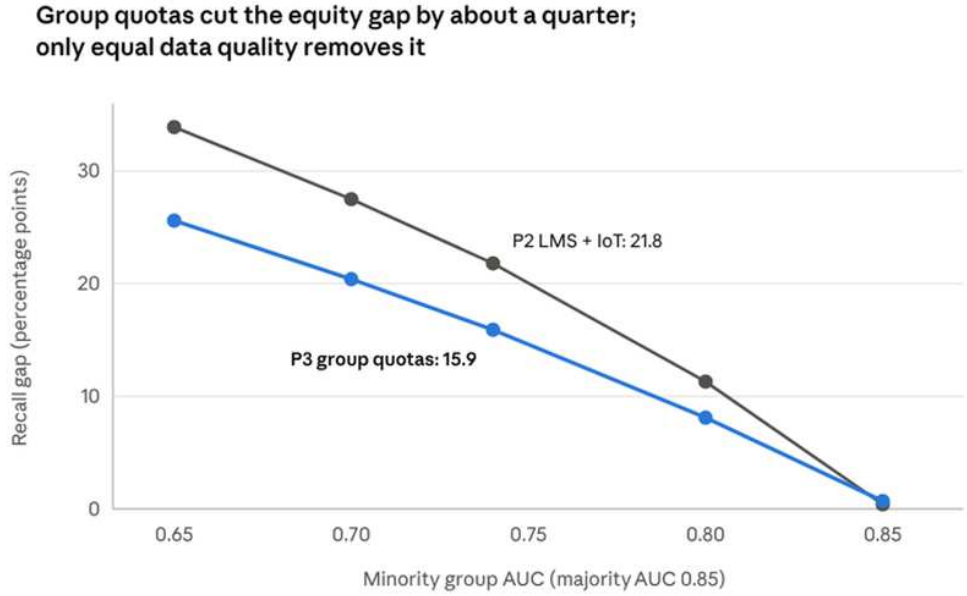


Figure 2. Equal-opportunity gap (majority minus minority recall) for P2 and P3 as the minority group's AUC varies (majority AUC 0.85; 300 replications per point).

At every minority AUC below 0.85, quotas reduced the gap by roughly a quarter (for example from 33.9 to 25.6 points at AUC 0.65) while overall recall fell by at most about half a percentage point (48.6% to 48.1% at AUC 0.65). In this paper with assumption, when both groups have equal data quality (AUC 0.85), the gap is near zero under both policies (0.4 and 0.7 points) within simulation noise. The implication is that quotas are a useful interim measure while the structural solution is better data for under-served groups, for instance, by closing connectivity or sensor-coverage gaps. This aligns with documented concerns that learning analytics should not assume one size fits all (Gašević et al., 2016) and that fairness depends on the criterion chosen (Hardt et al., 2016; Mehrabi et al., 2021).

IoT microclimate monitoring (RQ3)

A three-sensor consensus rule detected 98.0% of humidity excursions with 0.7 false alarms per 30 days while a fixed threshold detected 47.0% of excursions with 12.4 false alarms (Table 6).

Table 6. Experiment B baseline: alarm strategies (30 evaluation days, 4 excursions per replication; 300 replications).

Detector	Excursions detected	Mean delay among detected (min)	False alarms per 30 days
D1 fixed threshold (60% RH)	47.0%	109.0	12.35
D2 single-sensor EWMA	97.9%	79.3	2.12
D3 three-sensor consensus	98.0%	79.8	0.73

The fixed threshold sits above normal humidity peaks, so it ignores smaller excursions, and it also fires on single-sample spikes. The delay figure for D1 is not comparable with the others because it is computed only over the larger excursions that D1 detects. Figure 3 shows detection by excursion size.

ai_iot_edu_culture_sim.py · Experiment B · 300 replications

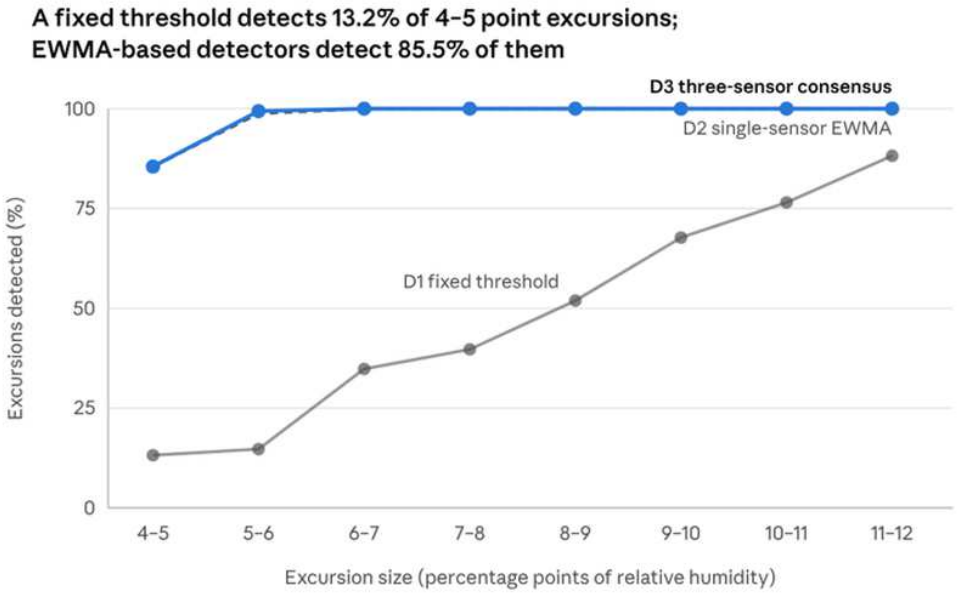


Figure 3. Detection rate by excursion size (percentage points of relative humidity above normal) for the three alarm strategies.

D1 detected only 13.2% of excursions of 4–5 points and 88.2% of those of 11–12 points, whereas the EWMA-based detectors identified 85.5% of the smallest excursions and 100% of those of six points or more. Preventive conservation requires timely responses to minimize the damaging effects on the treated material by such excursions (Sesana et al., 2021), which makes threshold a relatively weak control. Consensus eliminated most false alarms with only a marginal increase in response lag (79.8 versus 79.3 minutes).

Table 7. False alarms per 30 days as the sensor glitch rate varies (100 replications per row).

Glitch probability per sample per sensor	D1 fixed threshold	D2 single-sensor EWMA	D3 three-sensor consensus
0.0005	2.84	1.25	0.70
0.002	11.86	1.94	0.66
0.008	42.77	4.95	0.65

Table 7 demonstrates that false positives in the single-sensor detectors increase as the glitch rate increases, but the consensus rule generally maintains a low value near 0.7 alarms per 30 days at every level. Alarm fatigue is a real concern for small heritage teams, and the robustness of the consensus rule may be as valuable as its ability to detect threats. The absolute values depend on our choice of the glitch model, while true sensor faults might be correlated between sensors, reducing the benefit of consensus.

Governance framework (RQ4)

The results highlight four principles which cut across education and heritage, and which we suggest are critical in any such system: equity, privacy and security, cultural sovereignty, and human oversight. Figure 4 shows the cross-references in which, each principle to suggested practices, while Table 8 associates each documented risk with a control and an accountable agent.

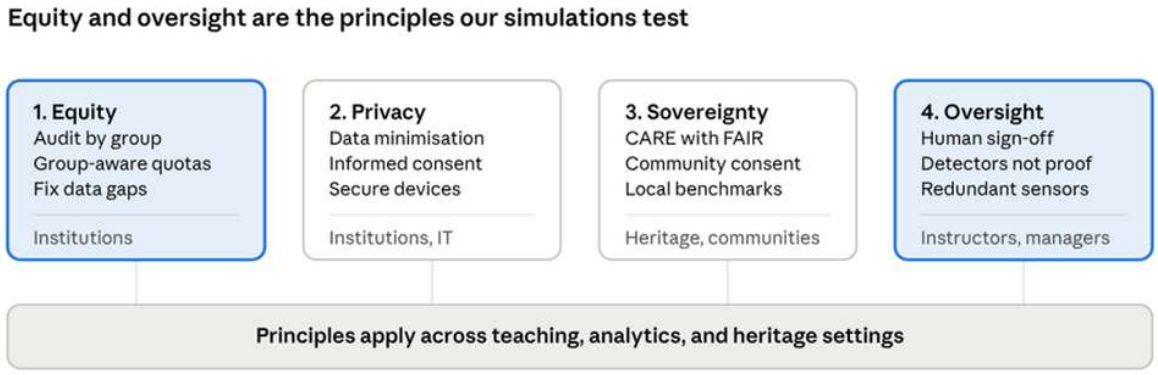


Figure 4. Four governance principles for AI and IoT in education and cultural studies, each with example practices.

Table 8. Risk–control matrix.

Risk	Evidence	Control	Primary actor
Unequal data quality across groups	Section 3.2; Baker & Hawn (2022)	Report recall by group before deployment; use group-aware allocation; invest in data quality for under-served groups	Institutions, analytics teams
Proxy-target bias	Obermeyer et al. (2019)	Validate the prediction target against the outcome that matters to students	Analytics teams
Privacy and consent in student and sensor data	Slade & Prinsloo (2013); Roman et al. (2013)	Data minimisation, transparent consent, secure device management	Institutions, IT services
Unreliable AI-text detection	Liang et al. (2023)	Do not use detector scores as sole evidence; rely on process records and dialogue	Instructors, institutions
Cultural value bias in LLM output	Tao et al. (2024)	Evaluate against locally relevant benchmarks and keep human review	Educators, developers
Alarm fatigue from sensor glitches	Section 3.3	Use redundant sensors and consensus rules	Facility and conservation managers
Data sovereignty over cultural material	Carroll et al. (2020); Wilkinson et al. (2016)	Apply CARE principles alongside FAIR; obtain community consent for sensitive material	Heritage institutions, communities

Discussion and limitations

Final Rewrite

The two experiments demonstrate one theme in two contexts. In Experiment A, a better signal (IoT features) increased average performance while widening a gap, because the additional data were less informative for the group that most needed attention. When a smarter detector was used in Experiment B, the advantage of redundancy outweighed its higher complexity per unit. In both cases, the deciding factor is the quality and reliability of the data pipeline, rather than the sophistication of the model. This is a similar to the general concern that AI in education can exacerbate existing inequities if the evaluation is average

(Baker & Hawn, 2022; Selwyn, 2022), and extended the cultural bias concerns expressed for LLMs (Tao et al., 2024) to the data pipeline.

Number of limitations apply to the work. First of all, all quantitative results are based on stated assumptions (Tables 2 and 3): that scores are binormal, minority AUC is assumed and not measured, prevalence is equal across groups, the intervention effect is a single assumed, and support capacity is fixed. Second one, Experiment B uses a synthetic humidity process, an assumed glitch model, and a fixed threshold, while real gallery climates, sensor faults, and conservation limits vary by the object and site. Third one, the synthesis of literature is purposive and not exhaustive, and the field is evolving fast. And fourth one, neither experiment includes real learners, instructors, or conservators, so acceptability, workload, and legal constraints are not modeled. An empirical validation using institutional data and pilot deployments would be the logical next steps.

4. Conclusion

AI and IoT, enrich in both education and cultural studies, but the benefit depends on the quality, coverage, and reliability of the data they base on, and the accountability of the people who make decisions. There are answers to the research questions, that are as follows:

- RQ1. In this paper purposive synthesis, evidence is richest for student analytics and AI-assisted teaching, IoT in education is largely descriptive, and cultural-studies applications have a thinner and more technical literature. The recurring risks are data quality inequality, inadequate privacy protection, unreliable AI-text detection, and encoded cultural values.

- RQ2. At the simulated cohort, random allocation achieved 15.2% of at-risk students, LMS analytics 37.4%, and LMS plus IoT data 50.2%. Adding IoT data increased majority recall by 15.1 points and minority recall by 3.9, increasing the equity gap from

10.6 to 21.7 points. Group quotas reduced the gap to 15.7 points with the same overall reach, and only equal data quality eliminated it.

- RQ3. A three-sensor consensus rule identified 98.0% of humidity excursions with 0.7 false alarms per 30 days, versus 47.0% detection and 12.4 false alarms for a fixed threshold.

- RQ4. In this paper all four principles follow: equity, privacy and security, cultural sovereignty, and human oversight, each one tied to controls and accountable actors as described in Table 8.

For institutions and practitioners, five recommendations are followed as:

1. Report group outcomes, not only averages, before deploying analytics, and apply group-aware allocation as an interim safeguard.

2. Invest in data quality and coverage for under-served groups, since added data can increase gaps if it is less informative to them.

3. Benchmark any early-warning model against random allocation under the same capacity, and state the intervention effect assumed.

4. In conservation settings, use redundant sensors with a consensus rule, and evaluate detectors on small, sustained excursions as well as large ones.

5. Treat privacy, consent, and cultural data sovereignty as design requirements, applying community governance principles alongside FAIR practice, and hold humans accountable for decisions that affect students and collections.

The future work should validate the models with real institutional and sensor data, estimate group-specific model performance in practice, test correlated sensor faults, evaluate interventions (rather than assume their effect), and explore how generative AI and IoT can interact in humanities teaching and museum learning.

5. References

1. Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A survey on enabling technologies, protocols, and applications. *IEEE Communications Surveys & Tutorials, 17*(4), 2347–2376. https://doi.org/10.1109/COMST.2015.2444095
2. Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. *Computer Networks, 54*(15), 2787–2805. https://doi.org/10.1016/j.comnet.2010.05.010
3. Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. *International Journal of Artificial Intelligence in Education, 32*, 1052–1092. https://doi.org/10.1007/s40593-021-00285-9

4. Bekele, M. K., Pierdicca, R., Frontoni, E., Malinverni, E. S., & Gain, J. (2018). A survey of augmented, virtual, and mixed reality for cultural heritage. *Journal on Computing and Cultural Heritage, 11*(2), Article 7. https://doi.org/10.1145/3145534
5. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of FAccT '21* (pp. 610–623). https://doi.org/10.1145/3442188.3445922
6. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology, 3*(2), 77–101. https://doi.org/10.1191/1478088706qp063oa
7. Carroll, S. R., Garba, I., Figueroa-Rodríguez, O. L., et al. (2020). The CARE Principles for Indigenous Data Governance. *Data Science Journal, 19*, 43. https://doi.org/10.5334/dsj-2020-043
8. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41*(3), Article 15. https://doi.org/10.1145/1541880.1541882
9. Chen, L., Chen, P., & Lin, Z. (2020). Artificial intelligence in education: A review. *IEEE Access, 8*, 75264–75278. https://doi.org/10.1109/ACCESS.2020.2988510
10. Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2024). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International, 61*(2), 228–239. https://doi.org/10.1080/14703297.2023.1148890
11. Ferguson, R. (2012). Learning analytics: Drivers, developments and challenges. *International Journal of Technology Enhanced Learning, 4*(5/6), 304–317. https://doi.org/10.1504/IJTEL.2012.051816
12. Fiorucci, M., Khoroshiltseva, M., Pontil, M., Traviglia, A., Del Bue, A., & James, S. (2020). Machine learning for cultural heritage: A survey. *Pattern Recognition Letters, 133*, 102–108. https://doi.org/10.1016/j.patrec.2020.02.017
13. Gašević, D., Dawson, S., Rogers, T., & Gasevic, D. (2016). Learning analytics should not promote one size fits all: The effects of instructional conditions in predicting academic success. *The Internet and Higher Education, 28*, 68–84. https://doi.org/10.1016/j.iheduc.2015.10.002
14. Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. *Future Generation Computer Systems, 29*(7), 1645–1660. https://doi.org/10.1016/j.future.2013.01.010
15. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems 29*. https://doi.org/10.48550/arXiv.1610.02413
16. Hwang, G.-J., Xie, H., Wah, B. W., & Gašević, D. (2020). Vision, challenges, roles and research issues of Artificial Intelligence in Education. *Computers and Education: Artificial Intelligence, 1*, 100001. https://doi.org/10.1016/j.caeai.2020.100001
17. Ji, Z., et al. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys, 55*(12), Article 248. https://doi.org/10.1145/3571730
18. Kasneci, E., et al. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences, 103*, 102274. https://doi.org/10.1016/j.lindif.2023.102274
19. Kassab, M., DeFranco, J., & Laplante, P. (2020). A systematic literature review on Internet of Things in education: Benefits and challenges. *Journal of Computer Assisted Learning, 36*(2), 115–127. https://doi.org/10.1111/jcal.12383

20. Lakkaraju, H., Aguiar, E., Shan, C., et al. (2015). A machine learning framework to identify students at risk of adverse academic outcomes. In *Proceedings of KDD '15* (pp. 1909–1918). https://doi.org/10.1145/2783258.2788620
21. Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. *Patterns, 4*(7), 100779. https://doi.org/10.1016/j.patter.2023.100779
22. Lucas, J. M., & Saccucci, M. S. (1990). Exponentially weighted moving average control schemes: Properties and enhancements. *Technometrics, 32*(1), 1–12. https://doi.org/10.1080/00401706.1990.10484583
23. Macfadyen, L. P., & Dawson, S. (2010). Mining LMS data to develop an "early warning system" for educators: A proof of concept. *Computers & Education, 54*(2), 588–599. https://doi.org/10.1016/j.compedu.2009.09.008
24. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys, 54*(6), Article 115. https://doi.org/10.1145/3457607
25. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187–192. https://doi.org/10.1126/science.adh2586
26. Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science, 366*(6464), 447–453. https://doi.org/10.1126/science.aax2342
27. Page, M. J., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ, 372*, n71. https://doi.org/10.1136/bmj.n71
28. Roman, R., Zhou, J., & Lopez, J. (2013). On the features and challenges of security and privacy in distributed Internet of Things. *Computer Networks, 57*(10), 2266–2279. https://doi.org/10.1016/j.comnet.2012.12.018
29. Romero, C., & Ventura, S. (2020). Educational data mining and learning analytics: An updated survey. *WIREs Data Mining and Knowledge Discovery, 10*(3), e1355. https://doi.org/10.1002/widm.1355
30. Selwyn, N. (2022). The future of AI and education: Some cautionary notes. *European Journal of Education, 57*(4), 620–631. https://doi.org/10.1111/ejed.12532
31. Sesana, E., Gagnon, A. S., Ciantelli, C., Cassar, J., & Hughes, J. J. (2021). Climate change impacts on cultural heritage: A literature review. *WIREs Climate Change, 12*(4), e710. https://doi.org/10.1002/wcc.710
32. Sicari, S., Rizzardi, A., Grieco, L. A., & Coen-Porisini, A. (2015). Security, privacy and trust in Internet of Things: The road ahead. *Computer Networks, 76*, 146–164. https://doi.org/10.1016/j.comnet.2014.11.008
33. Slade, S., & Prinsloo, P. (2013). Learning analytics: Ethical issues and dilemmas. *American Behavioral Scientist, 57*(10), 1510–1529. https://doi.org/10.1177/0002764213479366
34. Tao, Y., Viberg, O., Baker, R. S., & Kizilcec, R. F. (2024). Cultural bias and cultural alignment of large language models. *PNAS Nexus, 3*(9), pgae346. https://doi.org/10.1093/pnasnexus/pgae346
35. Wilkinson, M. D., et al. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data, 3*, 160018. https://doi.org/10.1038/sdata.2016.18
36. Williamson, B., & Eynon, R. (2020). Historical threads, missing links, and future directions in AI in education. *Learning, Media and Technology, 45*(3), 223–235. https://doi.org/10.1080/17439884.2020.1798995

37. Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education: Where are the educators? *International Journal of Educational Technology in Higher Education, 16*, 39. https://doi.org/10.1186/s41239-019-0171-0
38. Zhang, K., & Aslan, A. B. (2021). AI technologies for education: Recent research & future directions. *Computers and Education: Artificial Intelligence, 2*, 100025. https://doi.org/10.1016/j.caeai.2021.100025