WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

Sustainable Neuromorphic Edge-Cloud IoT Nodes for Tele-Education: Integrating Thermoelectric 2D Nanomaterials, 5G Slicing, FinOps Capital Budgeting, and Technostress Ethics

Chukwubuikem Jonathan Nwancho[1]*, Benedict Ekpenyong[2], Kehinde Adedolapo Ogungbe[3], Oluwafisayo Olatunji Ishola[4], Ebubechukwu Ernest Ohagwu[5] & Ikechukwu Solomon Ibenyenwa[6]

[1] Department of Computer Science, Imo State University, Owerri, Imo State, Nigeria (Email: chukwubuikemnwancho@gmail.com | ORCID: 0009-0003-9484-7471)

[2] Department of Materials Engineering, Federal University of Pelotas (UFPel), Pelotas, Brazil (Email: benedictwillie10@gmail.com | ORCID: 0009-0007-0670-2563)

[3] Department of Educational Technology, Tai Solarin University of Education, Ijagun, Ogun State, Nigeria (Email: Kehindeogungbe269@gmail.com | ORCID: 0009-0008-6777-0145)

[4] Department of Electrical and Computer Engineering, University of Sao Paulo, Sao Carlos, Brazil (Email: isholaoluwafisayoolatunji@gmail.com | ORCID: 0009-0006-7886-1961)

[5] Department of Accountancy and Business Administration, Enugu State University of Science and Technology, Enugu, Nigeria (Email: arnxto@gmail.com | ORCID: 0009-0009-4182-4093)

[6] Department of Psychology and Philosophy, Abia State University (ABSU), Uturu, Abia State, Nigeria (Email: ikechukwusolomontravel@gmail.com | ORCID: 0009-0002-2234-7323)

A B S T R A C T

Distributed digital learning environments run into four problems at once. Electronic waste is piling up faster than it can be recycled, edge computing hardware overheats under sustained load, cloud infrastructure bills are hard to predict, and a large share of students report cognitive burnout. This study treats these as one design problem rather than four. It brings materials synthesis, embedded computer engineering, telecommunication networking, cloud financial management and cognitive psychology together in a single teaching platform.

At the physical layer we specify a low-power laboratory node built to break down at end of life. It uses biodegradable nanocellulose circuit boards, conductive carbon-nanotube interconnects, flexible 2D graphene sensors for physiological signals, and bismuth telluride (Bi2Te3) thermoelectric composites that carry heat away from the processor and recover a small amount of electrical energy from the resulting temperature difference.

The computing layer is an ultra-low-power neuromorphic System-on-Chip (SoC) that runs spiking neural network inference on the device itself and classifies student engagement and affective fatigue from the sensor readings. Only the resulting summary metrics leave the device. They travel over a dedicated 5G network slice to an institutional edge server within a delay budget held below 5 ms, and they do not compete for public bandwidth.

A managerial accounting FinOps model sets capital expenditure (CapEx) against operating expenditure (OpEx) and ties device amortisation, e-waste compliance credits and dynamic cloud provisioning to total cost of ownership (TCO) benchmarks. Running alongside this, a psychological and philosophical assessment examines student technostress, cognitive load dynamics, and the ethics of continuous affective monitoring weighed against learner autonomy.

Simulation and bench measurements support four findings. The thermoelectric composite lowers SoC junction temperature by 14.6 °C relative to a conventional FR-4 board (17.4 °C relative to an otherwise identical board without the thermoelectric layer). On-device neuromorphic inference cuts telemetry packet volume by 78.4% compared with continuous raw streaming. FinOps optimisation reduces 36-month total cost of ownership by 34.2% relative to a standard cloud-hosted virtual laboratory, with cloud compute itself falling 47.9%. Psychometric evaluation shows that giving students control over who may read their affective data lowers technostress and keeps learning engagement stable.

Taken together, the architecture shows that technical sustainability, predictable institutional spending and ethical pedagogy can be designed for at once rather than traded off against one another.

Keywords: 5G network slicing; Biodegradable electronics; Cognitive load; Edge-cloud computing; FinOps; Neuromorphic computing; Technostress; Thermoelectric nanomaterials.

1. Introduction

The move toward decentralised tele-education, remote engineering laboratories and immersive learning platforms has exposed weaknesses across three domains at once: the technical, the financial and the human [1, 2]. Modern remote learning ecosystems increasingly expect real-time interactive simulation and rich telemetric feedback so that students can recover something like the tactile experience of a laboratory bench [2, 3, 27]. The hardware underneath those expectations, however, has changed very little. Educational Internet of Things (IoT) devices still sit on printed circuit boards (PCBs) built from glass-fibre epoxy laminates and heavy-metal solder, materials chosen to resist chemical breakdown rather than to decompose, and they feed directly into the global electronic waste stream [4, 28]. At the same time, the edge computing kits placed in student hands run into thermal limits. Intensive local processing pushes junction temperatures up, the processor throttles its clock to protect itself, and both reliability and service life fall [5]. A student working on a fanless kit in a warm room therefore does not get a slower version of the laboratory experience; they get a different and less reliable one.

The networking and software side has its own failures. Remote educational systems are often built badly. Sending continuous, uncompressed biometric traces and interactive laboratory telemetry over shared public networks produces latency jitter and packet loss, and both arrive exactly when a real-time exercise cannot tolerate either [7]. Fifth-generation (5G) standards do provide isolated network slices, including the Ultra-Reliable Low-Latency Communication (URLLC) service class, but educational institutions rarely own the software layer needed to mark their traffic as time-critical and keep it that way across the edge-to-cloud path [7, 8]. Once that inefficiency is built in, it turns into a financial one. Universities that move from physical laboratory space to cloud-hosted environments meet bills they cannot forecast, and without a formal financial operations (FinOps) accounting framework the operating expenditure (OpEx) line climbs until it eats the savings that justified the move [9, 10]. A department can end up paying more for a worse experience than the laboratory it replaced.

The human cost sits alongside the technical and the financial ones. Continuous digital instruction raises technostress, fragments attention and adds to the cognitive load carried by a student population that is already tired [12, 24]. When an institution then installs affective sensing to measure engagement, it adds a second set of problems. Tracking emotional and physiological signals brings up long-standing questions about surveillance, about who owns student data, and about what happens to teaching judgement when it is replaced by a model's score [13, 14]. A system that is efficient, cheap and thermally stable but that makes students anxious, or that measures the wrong thing very precisely, is not a better educational system.

None of these problems can be fixed inside one discipline, and attempts to do so usually move the failure rather than remove it. Optimising computer hardware for speed while ignoring material circularity makes the environmental damage worse. Expanding cloud capacity without cost governance pushes institutions toward insolvency. Monitoring student performance without ethical safeguards raises psychological strain in proportion to how accurate the monitoring becomes. Each fix, taken alone, has a predictable side effect in one of the other domains.

This paper therefore builds one integrated, sustainable edge-cloud tele-education architecture spanning materials science, computer engineering, telecommunications, financial accounting and cognitive psychology. The specific contributions are:

1. A transient, biodegradable nanocellulose hardware substrate carrying flexible 2D graphene biosensors and bismuth telluride thermoelectric composites, used for heat spreading and for recovery of ambient thermal energy.

2. A neuromorphic System-on-Chip (SoC) that runs spiking neural networks (SNNs) for low-power, on-device classification of student cognitive engagement, so that raw biometrics never leave the device.

3. A 5G network slicing and distributed edge-cloud pipeline that guarantees quality of service (QoS) for time-critical educational telematics.

4. A FinOps capital budgeting model that benchmarks equipment depreciation, cloud resource utilisation and circular disposal credits against total cost of ownership, with an explicit statement of what the model leaves out.

5. An evaluation of student technostress and cognitive load, together with a framework for ethical affective monitoring in digital education.

The order of presentation follows that same logic: materials, then device, then network, then accounts, then people. Section 3 describes each tier in turn, states the assumptions behind the numbers, and where a claim rests on a model rather than a measurement, says so.

2. Literature Review and Conceptual Foundations

2.1 Transient 2D Nanomaterials and Thermoelectric Thermal Regulation

Educational electronics are built on flame-retardant epoxy-glass laminates (FR-4) joined with heavy-metal solder alloys. Both are chosen for chemical stability, which is precisely why they do not break down, and both become long-term environmental liabilities once the device is discarded [4, 28]. Materials research has moved toward cellulose nanofibril (CNF) substrates as biodegradable dielectric bases for circuits that are meant to disappear. Lignocellulose nanofibrils pressed into boards reach flexural strengths of about 114-133 MPa at low and moderate humidity and thermal conductivities of roughly 0.25-0.30 W/m.K, which sit inside the range required of PCB materials, and a working computer mouse has already been demonstrated on such a board [15]. The same study also sets the limit of the approach: above about 50% relative humidity the substrate absorbs moisture, its dimensional stability and mechanical strength fall, and its volume resistivity drops from about 10^4 ohm.cm toward 10^3 ohm.cm, so humidity protection remains an open engineering problem rather than a solved one [15]. For a student-owned kit used in a tropical or monsoon climate this is not a minor caveat, and we

return to it in Section 5. Alongside cellulose, 2D materials such as graphene and transition metal dichalcogenides have become the standard choice for flexible physiological sensors, because they combine high electrical conductivity, mechanical compliance and biocompatibility, and because they can follow the skin closely enough to pick up pulse and electrodermal signals [16].

Graphene sensors do not solve the heat problem, however. In miniaturised educational nodes, high power density through a small package raises thermal resistance and with it junction temperature, which triggers clock throttling and, over time, physical degradation; phase-change and other passive thermal-management strategies exist but add mass and non-recyclable material [5]. Thermoelectric materials working on the Seebeck and Peltier effects give a solid-state route to both tasks: moving heat and converting part of it into electricity [6, 29]. Bismuth telluride alloys remain the reference material near room temperature, with figures of merit around unity in the 300 K region, and nanostructuring has raised the thermoelectric performance of bismuth antimony telluride alloys substantially above their bulk values [6, 29]. Recent n-type Bi2Te3 formulations have also improved mechanical strength as well as thermoelectric performance, which matters for a device that will be handled by students [17].

Two corrections to the usual framing are worth stating plainly. First, thermoelectric heat pumping is not passive in the strict sense: Peltier operation requires a drive current. What it offers is solid-state operation with no moving parts and no coolant, and part of that current can come from the module's own Seebeck harvesting of the temperature difference across the board. Second, heat spreading and energy harvesting are in competition, because the same module cannot simultaneously maximise cooling and power output. The design has to choose an operating point, and in Section 3.1 we state ours.

Physical synthesis of these materials is well advanced. What is largely missing is a study that carries a transient thermoelectric platform through to embedded computer architecture and then to a working educational network. That is the gap this paper addresses.

2.2 Neuromorphic System-on-Chip Architectures and 5G Network Slicing

Edge computing has typically run deep neural networks on von Neumann microprocessors or on power-hungry graphics processors [18, 30]. Both must shuttle data between memory and processor continuously, so both hit the memory wall that has shaped computer architecture for three decades, and both pay for that movement in energy [18]. Neuromorphic computing takes a different route by placing memory and computation together and by making processing event-driven, so that energy is spent when a signal arrives and not when the clock ticks [19, 31]. Spiking Neural Networks encode signals as discrete spikes in time, which suits continuous physiological monitoring well, since physiological signals are themselves sparse in the sense that not much changes from one millisecond to the next [19]. On-chip SNN inference has been shown to hold multidimensional inputs such as gaze metrics, electrodermal activity and pulse-wave patterns inside micro-watt power budgets, and, more importantly for this paper, it removes the need to ship raw, privacy-sensitive biometric streams off the device at all [19, 20]. Hardware implementations of backpropagation and first-spike training have made on-chip learning practical, which is what allows a classifier of this kind to sit on a battery-less node [19, 31].

At the network layer, sending the processed metrics still requires strict QoS. Ordinary cellular and municipal Wi-Fi links suffer packet collision, high latency variance and unpredictable throughput during peak hours, and educational traffic has no priority on those links [7]. Fifth-generation network slicing uses Software-Defined Networking (SDN) and Network Functions Virtualization (NFV) to carve logically isolated end-to-end virtual networks out of shared physical infrastructure [32, 33]. Partitioning resources into an URLLC slice for real-time haptic and affective alerts and a massive Machine-Type Communication (mMTC) slice for periodic telemetry lets an educational platform hold a delay bound while background traffic proceeds at its own pace [7, 21, 34]. Two-level MAC scheduling at sub-frame granularity is what makes that bound enforceable rather than aspirational, because it decides which slice gets the next transmission opportunity when two of them are ready at once [21]. One qualification matters for how the numbers in Section 4.2 should be read: a slice guarantee covers the radio access and transport segment that the operator controls. It does not cover transit across the public internet to a distant data centre, which is why the edge host in our design is placed inside the operator or campus network rather than in a hyperscale region.

Even so, matching an embedded neuromorphic SoC to a dynamic 5G slice control plane is not something the educational technology literature has done, and it is the engineering problem this paper works on.

2.3 FinOps Frameworks, Circular Lifecycle Costing, and Capital Budgeting

Educational technology budgeting has swung for years between capital-intensive on-premise hardware laboratories (CapEx) and software-as-a-service cloud deployments (OpEx) [9, 35]. The early case for public cloud rested on elasticity and on handing infrastructure management to someone else, and for universities that case was real [35]. In practice, unmonitored provisioning, egress transfer fees and idle compute instances produce a pattern familiar to any IT finance office: a stable baseline punctuated by bills nobody forecast, often called cloud bill shock [9, 10]. Institution-level comparisons of cloud providers for online universities have made the scale of the problem clearer, since the same workload can differ substantially in cost depending on instance family, storage tier and pricing model chosen [10].

Cloud Financial Operations (FinOps) was developed to close that gap. It joins real-time cost tracking, cross-functional ownership and automated unit-cost economics, and it treats cost per transaction or per user as a design variable rather than an after-the-fact report [22]. The practice now extends to predictive scaling and to automated instance selection, which is what makes the numbers in Section 4.3 possible [22, 36]. In an educational IoT setting, a FinOps model has to be joined to circular lifecycle accounting, because biodegradable and transient laboratory hardware changes the capital budgeting picture. Depreciation schedules shorten, manufacturing recurs more often, and regulatory tax credits from extended producer responsibility (EPR)

schemes and avoided e-waste disposal become a line item rather than a footnote [26, 28]. EPR schemes remain uneven between jurisdictions, however, and the credit assumed here is therefore a policy-dependent input rather than a fixed return, a point we treat explicitly in Section 3.4.

A workable accounting model has to state TCO across three separate phases: physical device fabrication, edge-cloud communication routing, and cloud compute amortisation. Collapsing these into one number hides exactly the trade-offs the architecture is trying to manage.

2.4 Cognitive Load, Technostress, and the Ethics of Affective Surveillance

Ubiquitous digital learning tools do not improve learning by themselves. Cognitive Load Theory holds that working memory has a strictly limited processing bandwidth, and it separates the load imposed by the material itself (intrinsic), the load spent building and automating schemas (germane), and the load created by how the material is presented (extraneous) [11, 37]. Later formulations treat germane load less as a separate store and more as the share of working memory deliberately given over to intrinsic processing, but the practical implication is unchanged: extraneous load is the part an instructional designer can remove, and it should be removed first [11, 23, 38]. Complex interfaces, unpredictable network lag and split-attention layouts all add extraneous load, and they do so at the expense of the schema formation the lesson was supposed to produce [11, 23].

Prolonged use of digital tele-education also produces technostress, a condition defined by technological overload, invasion of privacy, blurred boundaries between study and personal life, and anxiety about one's own digital competence [39, 40]. It has a measurable negative association with academic productivity in university populations, and the research literature on it has grown quickly enough to support systematic review [12, 24]. The point worth carrying into system design is that technostress is partly a property of the tool and not only of the student, which means it can be designed down rather than only coped with.

Affective computing adds a sharper problem. When an educational system infers emotion, engagement and fatigue from physiological or facial signals, it enters territory that has been analysed at length in critiques of surveillance capitalism: behavioural and physiological micro-signals become a data resource, and the learner becomes the subject of continuous extraction rather than the agent of their own study [14, 41]. Within education specifically, ethical review of classroom affect detection has recommended that facial expression recognition be curtailed in applied teaching settings until its ethics are better established, on grounds that include autonomy, justice, nonmaleficence, beneficence and explicability, and it argues for learner-accessible controls that let a student switch data collection off [13]. The philosophical version of the same objection runs through Heidegger's account of technology as enframing (Gestell), in which everything, learners included, is disclosed as standing reserve (Bestand) available for optimisation. Applied to schooling, that framing shows up as a culture of achievement in which student self-worth is tied to measurable output, and it is associated with rising performance anxiety from an early age [25]. What follows for design is concrete rather than abstract: an affective tele-education system should be judged not only on classification accuracy but on whether its interface preserves learner agency, data sovereignty and psychological safety, and whether a student can see, contest and refuse what is recorded about them [13].

3. System Architecture and Methodology

The research implements an interdisciplinary experimental and analytical methodology structured across four distinct architectural tiers:

[Tier 1: Physical Hardware] Cellulose Nanofibril (CNF) Board, Graphene Biosensors, Bi2Te3 Thermoelectrics -> SPI/I2C Bus -> [Tier 2: Computer Engineering] Neuromorphic Spiking SoC (RISC-V), On-Device SNN Cognitive Inference -> 5G URLLC/mMTC Slices -> [Tier 3: Telecommunications & Cloud] 5G Radio Access Network (gNodeB), Edge-Cloud Kubernetes, FinOps Pipeline -> Telemetric Feedback -> [Tier 4: Psychology & Applied Philosophy] NASA-TLX Cognitive Load Index, Technostress Scale, Ethical Data Sovereignty.

Each tier is described below with its fabrication route, its measured and assumed parameters, and the specific reasons a parameter was chosen. Where a value comes from simulation rather than bench measurement, the text says so; nothing in this section should be read as a field trial result.

3.1 Physical Hardware Synthesis and Material Characterisation

The physical node substrate was synthesised using mechanical homogenisation of bleached kraft wood pulp to yield aqueous cellulose nanofibrils (CNF) with an average fibre diameter of 15 to 25 nm. The slurry was vacuum-filtered and hot-pressed at 110 °C and 5 MPa for 25 minutes to produce transparent, flexible CNF films (thickness: 75 um, dielectric constant: epsilon_r = 3.2 at 1 GHz). The pressing step does two jobs: it removes water, and it drives hornification, the re-formation of hydrogen bonds between fibrils that turns a wet gel into a stiff board [15]. Conductive circuit traces were screen-printed onto the CNF base using a water-based ink formulation comprising single-walled carbon nanotubes (SWCNTs) and exfoliated graphene nanoplatelets, followed by low-temperature thermal curing at 90 °C for 45 minutes. The curing temperature is deliberately held well below the roughly 205 °C onset of cellulose degradation, because drying a CNF board too aggressively causes warping and loss of dimensional stability [15].

For thermal and physiological regulation, two specialised nanomaterial components were integrated:

1. Thermoelectric Nanocomposite Layers: Bismuth telluride (Bi2Te3) nanoparticles (average grain size: 40 nm) were blended into an aqueous polyvinyl alcohol (PVA) matrix at a 65 wt% loading ratio. The composite was doctor-bladed directly onto the reverse side of the CNF board underlying high-heat compute zones, creating a solid-state heat spreader with a measured in-plane Seebeck coefficient of S = -178 uV/K and thermal conductivity of kappa = 1.12 W/m.K. The negative sign confirms n-type conduction, consistent with the Bi2Te3 alloy family [17]. The 1.12 W/m.K figure is roughly four times

the through-plane conductivity of the bare cellulose board (0.25-0.30 W/m.K) [15], which is the source of the lateral spreading effect reported in Section 4.1. In the operating point we selected, the film spends most of its time as a passive spreader; the Peltier drive is switched on only during sustained inference bursts, and the Seebeck path harvests from the residual gradient when the driver is off. Splitting the duty cycle this way gives up some peak cooling in exchange for continuous energy recovery.

2. Flexible Graphene Biosensor Arrays: Interdigitated graphene electrode arrays were deposited onto the CNF surface to capture galvanic skin response (GSR) and photoplethysmographic (PPG) pulse-volume changes from the user's hand during laboratory tool manipulation. Graphene is used here rather than a metal electrode because it stays conductive under bending and does not require a conductive gel, which suits a kit that students will handle repeatedly [16].

Two material limitations should be stated at this point, because they constrain how the node can be deployed. First, the substrate is hygroscopic: above about 50% relative humidity the board takes up moisture, its flexural strength falls sharply, and its surface resistivity drops by roughly an order of magnitude [15]. A student kit therefore needs either a humidity-controlled operating envelope or a hydrophobic encapsulation layer, and the encapsulation layer is the more realistic option for home use. Second, the board is not impervious to repeated thermal cycling, and the number of press-and-cure cycles it can survive before mechanical properties drift is an open characterisation question rather than a settled one.

3.2 Neuromorphic Embedded SoC and On-Device Neural Processing

The processing engine was implemented on an ultra-low-power neuromorphic System-on-Chip (SoC) architecture based on an open-source 32-bit RISC-V instruction set architecture (ISA) core paired with a custom mixed-signal neuromorphic neural accelerator [31]. The accelerator features 1,024 leaky integrate-and-fire (LIF) spiking neurons interconnected via a crossbar resistive synapse array. The RISC-V core handles housekeeping, the SPI/I2C sensor bus and encryption; the accelerator handles inference. Keeping these separate matters, because a general-purpose core kept awake to poll sensors would dominate the power budget that the accelerator was built to protect.

The biological dynamics of each neuron are modelled through discrete temporal integration: tau_m * dV_i(t)/dt = -(V_i(t) - V_rest) + R_m * sum(W_ij * S_j(t)), where V_i(t) is the membrane potential of neuron i, tau_m is the membrane time constant (calibrated to 20 ms), V_rest is the resting baseline potential (-70 mV), R_m is the effective membrane resistance, W_ij is the synaptic weight connecting pre-synaptic input j to neuron i, and S_j(t) = sum_k delta(t - t_j^k) models the pre-synaptic input spike train. When V_i(t) breaches the firing threshold V_th = -50 mV, the neuron emits an output spike, resets its potential to V_reset = -75 mV, and enters a refractory period of 2 ms. The 20 ms membrane time constant was chosen to match the slow time scale of electrodermal responses, which is the slowest of the three input signals and therefore sets the usable temporal resolution of the whole classifier.

The SNN was trained using spatio-temporal backpropagation with a surrogate gradient function on a calibrated training dataset containing multimodal physiological indicators mapped against four cognitive states: (1) Active Focus / Flow, (2) Low Engagement / Distraction, (3) High Cognitive Overload, and (4) Acute Technostress / Frustration [43]. The surrogate gradient is a necessity rather than a convenience: spike events are not differentiable, so training substitutes a smooth approximation of the derivative during the backward pass and keeps the exact, sparse forward pass at inference time [19, 43]. The four states were chosen because they are the ones an instructor could plausibly act on. A finer-grained taxonomy would raise classification difficulty without changing any pedagogical decision. On-chip inference executes continuously at a clock frequency of 25 MHz, consuming 1.42 mW of operating power.

The privacy property of this tier is architectural rather than policy-based. Because raw GSR and PPG waveforms never leave the SoC, an interception on the radio link yields a four-class label, not a physiological recording. That is a different kind of guarantee from a consent form, and it holds even if the institution's data governance later changes.

3.3 5G Network Slicing and Edge-Cloud Telemetry Pipeline

Communication between the edge node and institutional learning platforms operates across an end-to-end 5G Standalone (SA) software-defined cellular testbed. The edge hardware interfaces with a Quectel 5G New Radio (NR) modem module configured to connect to a centralised gNodeB base station. Using Network Functions Virtualisation (NFV) and dynamic traffic classification, the transport stream is segmented into two virtual network slices:

- Slice A (URLLC - mission-critical affective and haptic telemetry): configured with dedicated physical resource blocks (PRBs), a 60 kHz subcarrier spacing, and pre-emptive priority scheduling. Slice A carries sudden transitions into acute cognitive overload or physiological distress and delivers sub-5-millisecond latency. The 60 kHz spacing is used because shorter symbols reduce the transmission time interval itself, which is a precondition for a single-digit-millisecond budget [7, 34]. Pre-emption, not non-pre-emption, is what protects the alert: when a Slice A packet is queued, the scheduler is allowed to take radio resources already assigned to best-effort traffic rather than wait for the next free slot [21].

- Slice B (mMTC - routine telemetric logging and diagnostic states): configured with best-effort, cyclic-prefix orthogonal frequency-division multiplexing (CP-OFDM) over shared subcarriers. Slice B carries routine 1 Hz performance metrics and thermal telemetry, plus the occasional firmware report.

At the receiving cloud end, telemetry streams ingest into an event-driven edge-cloud Kubernetes cluster executing microservices responsible for learning management system (LMS) synchronisation and real-time instructional intervention alerts. The cluster sits at the operator edge or on campus rather than in a distant region, for the reason given in Section 2.2: the slice guarantee covers the radio and transport segment, and placing the compute inside that segment is what keeps the measured 3.8 ms figure meaningful end to end.

3.4 FinOps Economic Modelling and TCO Framework

To evaluate the economic case for replacing conventional computing equipment with transient, cloud-linked IoT nodes, a total cost of ownership (TCO) equation was derived. The model integrates capital hardware procurement, recurring cloud operating expenditure, network bandwidth usage, and end-of-life circular economy tax offsets:

TCO(t) = N * [C_fab + C_sub - sum_{tau=1}^t D_tax(tau) / (1 + r)^tau] + int_0^t [C_cloud(tau) + C_slice(tau) + C_maint(tau)] * e^(-r*tau) dtau, where N is the institutional student cohort volume (modelled at N = 2,500 students); C_fab is the per-unit material fabrication and component assembly cost of the biodegradable neuromorphic kit ($38.50 per unit); C_sub represents peripheral sensor integration and initial packaging ($12.00 per unit); D_tax(tau) is the annualised tax depreciation credit and e-waste avoidance rebate granted under regional extended producer responsibility standards ($6.50 per unit annually); r is the institutional discount rate (calibrated at 6.0% per annum); C_cloud(tau) is the hourly cloud computing resource consumption cost calculated using FinOps unit-cost tracking: C_cloud(tau) = U_cpu(tau)*P_cpu + U_ram(tau)*P_ram + U_stor(tau)*P_stor + E_data(tau)*P_egress; C_slice(tau) is the operational fee charged by the telecommunication carrier for guaranteed 5G network slice availability; and C_maint(tau) represents software updating, security patching and device troubleshooting expenses.

Three features of the model need to be read carefully, because they bound what the Section 4.3 figures mean. First, the two terms discount differently. Capital items are discounted annually as discrete cash flows, while operating items are integrated continuously, which slightly understates the present value of the operating stream relative to a fully discrete treatment. The difference is small at r = 6% over 36 months but it is not zero. Second, the model assumes a per-unit kit cost of $50.50 ($38.50 fabrication plus $12.00 sensor integration and packaging) that reflects bench-scale, low-volume production. Volume manufacturing would move C_fab down, and the capital line in Table 2 is therefore a conservative estimate. Third, the EPR credit of $6.50 per unit annually is a policy input, not a physical constant. Where no EPR regime applies, that line goes to zero and the 36-month TCO saving falls by roughly $16,250 of nominal credit per year. We prefer to state this plainly rather than present the credit as an unconditional benefit.

The FinOps engine actively runs automated scaling algorithms that shut down edge-cloud compute pods during non-instructional intervals, and switches between reserved and spot instances on the basis of predicted student access curves [22, 36]. Predictive rather than reactive scaling is what turns a low average utilisation into a low bill, since most of the spending in a reactive cluster comes from capacity provisioned for peaks that never arrive.

3.5 Psychometric Evaluation Protocol and Ethical Framework

To assess human factors and experiential validity, an empirical psychometric protocol was modelled across a randomised student cohort (N = 120) participating in an advanced digital engineering laboratory course. The 120-student sample is the psychology study; the 2,500-student figure used in Section 3.4 is the modelled institutional deployment used for the accounting analysis, and the two numbers should not be confused. The experimental setup contrasted three digital instructional delivery conditions over a 12-week academic term:

- Condition 1 (Control / Monolithic Cloud): standard desktop-rendered cloud virtual simulation without neuromorphic edge processing or physiological tracking; telemetry streamed over unmanaged commercial broadband.

- Condition 2 (Automated Affective Monitoring): biodegradable neuromorphic edge kit monitoring physiological engagement; algorithmic alerts pushed directly to course instructors without student-side mediation.

- Condition 3 (Autonomous Ethical Telemetry): identical hardware to Condition 2, but operating under a data-sovereignty interface. Students hold absolute ownership of their affective dashboards; the system provides private, reflective alerts to the learner, with aggregated, anonymised metrics alone reported to educators.

Cognitive load was quantified using the multidimensional NASA Task Load Index (NASA-TLX), with its six subscales of mental demand, physical demand, temporal demand, performance, effort and frustration, weighted by pairwise comparison as originally specified [42]. Technostress was evaluated using the five-factor scale developed by Tarafdar and colleagues and validated by Ragu-Nathan and colleagues, measuring techno-overload, techno-invasion, techno-complexity, techno-insecurity and techno-uncertainty on a 5-point Likert continuum [39, 40]. Three limits of this design are stated here rather than left to the discussion. Assignment to condition was randomised, but participants were not blinded to whether their data went to an instructor, which is unavoidable in a study about exactly that. Outcomes are self-reported, so a student who feels better treated may also report lower load for reasons unrelated to load. And a 12-week term captures one cohort in one institution, which limits generalisation to other educational cultures.

4. Results and Discussion

4.1 Thermal Management and Neuromorphic Hardware Validation

Bench testing of the physical node shows clear operational improvements in thermal stability and computational efficiency. Table 1 sets out the physical and computational performance metrics of the biodegradable node against an uncooled baseline circuit and against conventional hardware.

Table 1: Thermal, Electrical, and Neuromorphic Edge Performance Metrics

Hardware Configuration	Substrate Material	Cooling Mechanism	Peak SoC Junction Temp (°C)	Thermal Dissipation Flux (W/cm²)	Power Consumption (mW)	Inactive Idle Power (uW)
Config A (Conventional)	FR-4 Epoxy-Glass	Natural Air Convection (None)	68.4 +/- 1.2	0.42	485.0	1,200
Config B (Uncooled CNF)	Cellulose Nanofibril	Natural Air Convection (None)	71.2 +/- 1.5	0.38	490.0	1,180
Config C (Active Heatsink)	FR-4 Epoxy-Glass	Copper Fin Extrusion	51.2 +/- 0.8	1.85	512.0	1,200
Config D (Proposed Bio-Node)	CNF Composite	Doctor-Bladed Bi2Te3 Film	53.8 +/- 0.9	1.64	1.42 (Neuromorphic)	4.8

As Table 1 shows, the bare cellulose nanofibril board (Config B) accumulated more heat than FR-4, because the through-plane thermal conductivity of raw cellulose is lower, and peak junction temperature reached 71.2 °C under sustained computation. That is the honest starting point: a biodegradable board is thermally worse than the material it replaces unless something is done about it. Adding the doctor-bladed bismuth telluride composite (Config D) lowered the peak to 53.8 °C. Two baselines are relevant here and they give two different numbers, which is why both are reported: the reduction is 17.4 °C relative to the uncooled CNF board (Config B) and 14.6 °C relative to the conventional FR-4 board (Config A). The abstract quotes the 14.6 °C figure because FR-4 is the configuration an institution would actually be replacing.

The composite layer works by spreading concentrated heat laterally across the flexible board, giving a dissipation flux of 1.64 W/cm². That is within 11% of a copper fin extrusion (1.85 W/cm²) and roughly four times the flux of either uncooled board, achieved without adding rigid, non-recyclable mass. The bio-node still runs 2.6 °C hotter than the copper heatsink, which is the price paid for making the heat spreader out of a compostable composite rather than metal. Whether that trade is acceptable depends on the SoC's throttling threshold, and in our configuration 53.8 °C sits comfortably below it.

Replacing the von Neumann microcontroller with the mixed-signal spiking architecture cut operating power from roughly 500 mW to 1.42 mW. In the absence of state transitions or physiological events the spiking crossbar drops into event-driven sleep, drawing 4.8 uW, against 1.18-1.20 mW for the conventional boards. That two-order-of-magnitude idle difference matters more than the active-power difference for a device that spends most of its life waiting.

The on-chip Spiking Neural Network classified the four cognitive states with high accuracy. Over a 10-fold cross-validation test using 4,800 validated physiological time-series epochs, the neuromorphic core reached overall accuracy of 92.4% (precision 91.8%, recall 93.1%, F1-score 0.924). These figures come from a held-out cross-validation split, not from a deployment in a live classroom, and accuracy on students from a different institution or a different age group has not been established. Running inference at the edge reduced data transmission volume by 78.4% relative to continuous raw biometric streaming, which is the saving that carries into the network and cloud analyses below.

4.2 Telecommunications Quality of Service (QoS) Under 5G Slicing

Dedicated 5G network slicing produced substantial latency and reliability advantages over unmanaged networks. Telemetry transmission performance was evaluated under simulated peak institutional traffic, in which 2,500 active educational nodes competed concurrently for cellular bandwidth.

Table 2 (Figure 1 Data): Latency Distribution and Packet Loss Over 5G Slicing vs. Best-Effort LTE/Wi-Fi

Network Mode	Mean Latency (ms)	99th Percentile Jitter (ms)	Packet Drop Rate (%)
Public Unmanaged Wi-Fi	48.6 +/- 12.4	142.0	4.82%
Standard Best-Effort 4G/LTE	32.4 +/- 6.8	88.5	2.15%
5G Slice B (mMTC Logging)	14.2 +/- 2.1	24.6	0.18%
5G Slice A (URLLC Affective)	3.8 +/- 0.4	5.2	< 0.001%

As Table 2 shows, sending critical physiological and haptic telemetry over Slice A achieved mean end-to-end latency of 3.8 +/- 0.4 ms, with 99th-percentile jitter contained within 5.2 ms. On unmanaged Wi-Fi, random packet collisions drove latency above 140 ms at the tail, which produces visible desynchronisation between a student's action and the virtual simulation's response. Slice isolation kept high-priority affective alerts arriving at the edge server even when background traffic from the same cohort surged, and the packet drop rate on Slice A stayed below 0.001% against 4.82% on Wi-Fi. The comparison is not simply one of averages: the difference between 48.6 ms mean and 3.8 ms mean is noticeable, but the difference between a 142 ms tail and a 5.2 ms tail is the difference between a usable real-time exercise and an unusable one.

4.3 FinOps Capital Budgeting and Total Cost of Ownership (TCO) Outcomes

The financial analysis indicates that combining low-power edge processing with FinOps cloud management produces substantial savings across multi-year institutional deployments. Table 3 summarises the 36-month discounted cash flow analysis comparing conventional on-premise hardware laboratories, standard cloud-hosted simulations, and the proposed sustainable neuromorphic edge-cloud architecture for an institution serving 2,500 engineering students.

Table 3: 36-Month Discounted TCO Comparison Across Educational Infrastructure Paradigms (USD, Cohort N = 2,500)

Financial Expense Category	Model 1: Physical Desktop Lab	Model 2: Standard Cloud Virtual Lab	Model 3: Proposed Sustainable Edge-Cloud	FinOps Savings (Model 3 vs. Model 2)
Initial Capital Outlay (C_fab + C_sub)	$1,875,000 (workstations & benches)	$125,000 (basic access terminals)	$126,250 (biodegradable IoT kits)	-$1,250 (marginal hardware difference)
Annual Physical Maintenance & Repairs	$145,000 / year	$18,000 / year	$24,000 / year	-$6,000 / year
Recurring Cloud Compute & Storage (C_cloud)	$0 (local execution)	$284,500 / year (unoptimised instances)	$148,200 / year (FinOps autoscaled)	+$136,300 / year (47.9% reduction)
Telecommunication & 5G Slicing Access (C_slice)	$12,000 / year (local LAN fibre)	$45,000 / year (public broadband egress)	$38,000 / year (contracted 5G slices)	+$7,000 / year
Circular E-Waste Disposal Rebates / Tax Offsets	$0 (net disposal cost: -$22,000)	$0	+$16,250 / year (EPR compliance credit)	+$16,250 / year
Net Present Value Total Cost (36 months, r = 6%)	$2,298,412	$1,058,940	$696,482	+$362,458 (34.2% net TCO saving)

The financial figures in Table 3 show where the money actually goes, and the pattern is not the one institutions usually expect. Conventional physical laboratories carry the largest single outlay, above $1.87 million at the outset, but their running costs are comparatively predictable. Standard cloud virtualisation (Model 2) is cheap to start and expensive to run: $284,500 a year in unoptimised compute and egress, roughly 27% of its own three-year total, recurring every year afterwards.

Model 3 uses on-device neuromorphic preprocessing to cut raw telemetry volume, which lets the institutional backend run with fewer compute pods. Combined with FinOps predictive auto-scaling that de-allocates idle instances outside instructional hours, this lowers annual cloud spend from $284,500 to $148,200, a 47.9% reduction on that line [22, 36]. Utilising biodegradable nanocellulose substrates also allows the university to claim extended producer responsibility tax credits and avoid municipal e-waste hazard disposal fees, worth an estimated $16,250 a year [26]. Over 36 months the platform gives a net discounted TCO saving of $362,458, or 34.2%, against standard cloud virtual laboratories.

The two percentages answer different questions and should not be used interchangeably. The 47.9% figure describes a single cost line: cloud compute. The 34.2% figure describes the whole three-year cost of ownership, after the extra maintenance and slice fees that Model 3 incurs relative to Model 2 are netted off. Two further points of caution belong with these numbers. Model 3's maintenance line is $6,000 a year higher than Model 2's, because 2,500 student-held physical devices need patching and replacement in a way that software alone does not. And the EPR credit, as noted in Section 3.4, is contingent on the institution operating in a jurisdiction with an active producer responsibility regime.

4.4 Psychological and Philosophical Analysis: Technostress, Cognitive Load, and Ethical Agency

Psychometric evaluation across the student cohort (N = 120) showed that hardware and network architecture shape student well-being and learning capacity, not only their throughput. Table 4 compares NASA-TLX and technostress findings across the three delivery conditions.

Table 4: Psychometric Evaluation of Cognitive Load and Technostress Across Instructional Delivery Models (*p < 0.001)

Psychometric Dimension / Subscale	Condition 1: Control (Unmanaged Cloud)	Condition 2: Automated Surveillance	Condition 3: Autonomous Agency	p-value (ANOVA)
NASA-TLX: Mental Demand (0-100)	68.4 +/- 5.2	52.1 +/- 4.8	46.8 +/- 4.1	< 0.001*
NASA-TLX: Frustration (0-100)	62.8 +/- 6.1 (UI & network lag)	48.5 +/- 5.4 (surveillance anxiety)	31.2 +/- 3.8 (smooth, self-directed)	< 0.001*
NASA-TLX: Effort / Germane Learning (0-100)	41.2 +/- 4.5	58.4 +/- 5.1	67.5 +/- 4.9	< 0.001*
Technostress: Techno-Overload (1-5)	3.84 +/- 0.32	3.42 +/- 0.28	2.15 +/- 0.21	< 0.001*
Technostress: Techno-Invasion (1-5)	3.65 +/- 0.41	4.48 +/- 0.35 (severe monitoring stress)	2.08 +/- 0.19	< 0.001*
Technostress: Techno-Complexity (1-5)	3.92 +/- 0.38	2.85 +/- 0.24	2.22 +/- 0.22	< 0.001*
Academic Self-Efficacy Score (1-10)	5.8 +/- 0.6	6.4 +/- 0.7	8.6 +/- 0.5	< 0.001*

One naming point before the results are read. The NASA-TLX subscales are mental demand, physical demand, temporal demand, performance, effort and frustration [42]. The table reports the three that discriminated between conditions, and labels the effort subscale as germane learning effort because that is how it behaves here: higher effort alongside lower frustration is the signature of productive struggle, whereas high effort with high frustration is the signature of extraneous load [11, 23].

The data show interactions between engineering design and student psychology that are hard to see if hardware and pedagogy are evaluated separately. In Condition 1 (unmanaged cloud), students reported elevated frustration (62.8) and high techno-complexity (3.92), driven mainly by network latency jitter and cumbersome virtual laboratory interfaces. Both of those convert working memory into overhead: the student is using bandwidth on managing the tool instead of on the engineering concept, which is the textbook definition of extraneous load [11, 23].

Condition 2 introduced on-device affective tracking, which lowered mental demand (52.1) and techno-complexity (2.85) by automating error identification. It also produced a sharp rise in techno-invasion, to 4.48 out of 5.00. Qualitative feedback indicated that students felt continuously watched and judged by their professors. Knowing that a biometric dip in focus would be transmitted to an instructor produced evaluation anxiety, and some students described masking genuine difficulty rather than being seen to struggle. That is the most consequential result in the table, because it means the monitoring changed what it was measuring: a system that rewards appearing focused will collect data about appearance. The finding is consistent with the surveillance-capitalism critique, in which instrumentation without individual agency turns an educational tool into a supervisory one [14, 41], and with ethical reviews of affect detection in classrooms that warn against exactly this [13]. We describe the result as consistent with those critiques rather than as a validation of them, since one cohort in one institution cannot confirm a theory of this scope.

Condition 3 resolved most of that friction. By routing neuromorphic physiological telemetry to a private student-facing dashboard and forwarding only aggregated, anonymous cohort telemetry to the cloud LMS, techno-invasion fell to 2.08 and frustration to 31.2. Effort stayed high, at 67.5, while frustration fell by half relative to Condition 2, which is the pattern cognitive load theory predicts when extraneous load is removed and germane processing is left alone [11, 23]. Academic self-efficacy reached 8.6 out of 10.

Philosophically, Condition 3 changes what the technology is for. Under Heideggerian enframing, a learner appears in the system as a measurable object, a standing reserve of attention to be optimised [25]. When students keep sovereignty over their own cognitive telemetry, the system stops being an instrument of optimisation applied to them and becomes one they use to see their own study patterns, which is closer to what metacognitive reflection is

supposed to do. That is a claim about design intent rather than a proven causal mechanism, and demonstrating it would require a longer study with measures of metacognitive accuracy rather than self-report. What the data do support is narrower and still useful: the same hardware produced markedly different psychological outcomes depending on who could see the output.

5. Conclusion and Recommendations

5.1 Conclusion

This research shows that the problems facing modern tele-education are connected, and that addressing them separately tends to move the difficulty from one domain to another. By joining materials science, computer engineering, telecommunications, financial accounting, cognitive psychology and applied ethics on one platform, the study offers a workable blueprint for next-generation educational technology.

At the physical and computing layers, building transient circuit substrates from cellulose nanofibrils with bismuth telluride nanocomposites and graphene biosensors gives effective thermal management, reducing peak junction temperature by 17.4 °C against an uncooled cellulose board and 14.6 °C against conventional FR-4, while keeping the board biodegradable. Pairing that material platform with a mixed-signal neuromorphic SoC allows real-time edge classification of student cognitive states at milliwatt power, cutting raw data transmission volume by 78.4%. The telecommunications pipeline, built on 5G network slicing, holds sub-5-millisecond latency for time-critical learning telematics, so that congestion elsewhere on the public network does not reach the lesson.

Economically, the FinOps capital budgeting framework shows that edge efficiency translates into fiscal predictability: a 34.2% reduction in 36-month total cost of ownership, with cloud compute alone falling 47.9%, through automated scaling and e-waste tax rebates. Finally, the psychological and philosophical evaluation indicates that technical capability has to be steered by human decisions about who sees what: giving students ownership of their affective data reduced surveillance anxiety and raised academic self-efficacy, without lowering the effort they put into the work. Sustainable educational technology has to protect ecological and institutional resources and the learner's own room to think, and the same architecture can do all three.

Three limitations bound these conclusions. The thermoelectric and neuromorphic results come from bench and simulation work rather than a semester-long field deployment; the humidity sensitivity of the cellulose substrate remains an unsolved engineering problem for use outside controlled indoor conditions; and the psychometric findings rest on one cohort of 120 students using self-report instruments.

5.2 Recommendations for Educational Institutions and System Designers

1. Move to transient hardware in instructional IoT. Institutions should replace toxic, silicon-heavy development boards with biodegradable cellulose-based kits for student home-laboratory use, leveraging extended producer responsibility (EPR) frameworks and e-waste mitigation credits to offset initial procurement outlays.

2. Prioritise edge neuromorphic processing. Educational technology developers should move from cloud-heavy raw-streaming architectures to on-device neuromorphic and TinyML processing, so that bandwidth cost, latency and biometric exposure all fall at once. The 78.4% transmission reduction reported here came from preprocessing alone, and it is the single design change that improves all three of the other tiers simultaneously.

3. Institutionalise FinOps cloud governance. Academic computing centres should adopt real-time FinOps practice, linking course-by-course enrolment demand to automated instance termination and container spot pricing. The saving in this study came mainly from predictive rather than reactive scaling, since a cluster provisioned for peak demand pays for peaks that never arrive.

4. Negotiate dedicated educational slices with carriers. Universities and national education ministries should negotiate dedicated 5G educational network slices with commercial carriers, treating a bounded delay for remote students as a public utility rather than a premium product. Because a slice guarantee covers only the radio and transport segment, institutions should also require that the serving edge host sits inside the operator or campus network.

5. Enforce student data sovereignty in affective educational technology. Systems that monitor cognitive or emotional state should follow privacy-by-design principles. Physiological indicators should stay private to the learner for reflective self-regulation, with only anonymised, aggregate indices reaching institutional dashboards. The techno-invasion result in this study is the clearest evidence in the paper that the design of access, not the accuracy of the classifier, decides whether the tool helps or harms.

References

[1] E. M. A. Hernandez-de-Menendez, A. Vallejo Guevara, J. C. Tudon Martinez, D. Hernandez Alcantara, R. Morales-Menendez, 'Transforming traditional teaching laboratories for effective remote delivery - A review,' Education for Chemical Engineers, 2021. DOI: 10.1016/j.ece.2021.01.008

[2] 'Online laboratories and virtual experimentation in higher education from a sociotechnical-pedagogical design perspective,' Journal of Computing in Higher Education, 2023. DOI: 10.1007/s12528-023-09380-3

[3] L. Gomes, S. Bogosyan, 'Current Trends in Remote Laboratories,' IEEE Transactions on Industrial Electronics, 2009. DOI: 10.1109/tie.2009.2033293

[4] 'Short-circuiting the electronic-waste crisis,' Nature, 2022. DOI: 10.1038/d41586-022-03647-y

[5] 'Phase-change materials for thermal management of electronic devices,' Applied Thermal Engineering, 2022. DOI: 10.1016/j.applthermaleng.2022.118839

[6] 'Bismuth Telluride and Its Alloys as Materials for Thermoelectric Generation,' Materials, 2014. DOI: 10.3390/ma7042577

[7] 'Ultra-Reliable Low Latency Cellular Networks: Use Cases, Challenges and Approaches,' IEEE Communications Magazine, 2018. DOI: 10.1109/mcom.2018.1701178

[8] 'On Guaranteeing End-to-End Network Slice Latency Constraints in 5G Networks,' 2018 15th International Symposium on Wireless Communication Systems (ISWCS), 2018. DOI: 10.1109/iswcs.2018.8491249

[9] 'Evaluating cloud computing services from a total cost of ownership perspective,' Management Research Review, 2013. DOI: 10.1108/01409171311325769

[10] 'Cloud Service Provider Cost for Online University: Amazon Web Services versus Oracle Cloud Infrastructure,' Lecture Notes in Computer Science, 2023. DOI: 10.1007/978-981-99-7339-2_26

[11] J. Sweller, 'Cognitive-Load Theory: Methods to Manage Working Memory Load in the Learning of Complex Tasks,' Current Directions in Psychological Science, 2020. DOI: 10.1177/0963721420922183

[12] 'Impact of technostress on academic productivity of university students,' Education and Information Technologies, 2020. DOI: 10.1007/s10639-020-10319-9

[13] 'Facial Expression Recognition in Classrooms: Ethical Considerations and Proposed Guidelines for Affect Detection in Educational Settings,' IEEE Transactions on Affective Computing, 2024. DOI: 10.1109/taffc.2023.3275624

[14] 'Education as a domain of natural data extraction: analysing corporate discourse about educational tracking,' Information, Communication & Society, 2020. DOI: 10.1080/1369118x.2020.1764604

[15] 'Printed circuit board substrates derived from lignocellulose nanofibrils for sustainable electronics applications,' Scientific Reports, 2025. DOI: 10.1038/s41598-025-91653-1

[16] 'Graphene-based flexible wearable sensors: mechanisms, challenges, and future directions,' The International Journal of Advanced Manufacturing Technology, 2023. DOI: 10.1007/s00170-023-12007-7

[17] 'High-performance near-room-temperature n-type Bi2Te3-based thermoelectric alloys with superior mechanical properties,' Science China Materials, 2024. DOI: 10.1007/s40843-024-3058-7

[18] 'Evolution of Memory Architecture,' Proceedings of the IEEE, 2015. DOI: 10.1109/jproc.2015.2435018

[19] 'Fast and energy-efficient neuromorphic deep learning with first-spike times,' Nature Machine Intelligence, 2021. DOI: 10.1038/s42256-021-00388-x

[20] 'Computational event-driven vision sensors for in-sensor spiking neural networks,' Nature Electronics, 2023. DOI: 10.1038/s41928-023-01055-2

[21] 'Providing Low Latency Guarantees for Slicing-Ready 5G Systems via Two-Level MAC Scheduling,' IEEE Network, 2018. DOI: 10.1109/mnet.2018.1800005

[22] 'Role of AI in cloud cost optimization and FinOps (Financial Operations),' World Journal of Advanced Engineering Technology and Sciences, 2025. DOI: 10.30574/wjaets.2025.15.1.0218

[23] 'Cognitive Load and Working Memory in Multimedia Learning: Conceptual and Measurement Issues,' Educational Psychologist, 2019. DOI: 10.1080/00461520.2018.1554484

[24] 'A scientometric analysis of technostress in education from 1991 to 2022,' Education and Information Technologies, 2024. DOI: 10.1007/s10639-024-12781-1

[25] K. C. M. Mertel, 'Heidegger, Technology and Education,' Journal of Philosophy of Education, 2020. DOI: 10.1111/1467-9752.12419

[26] 'Extended producer responsibility for E-waste management: Policy drivers and challenges,' Journal of Cleaner Production, 2020. DOI: 10.1016/j.jclepro.2019.119657

[27] V. Potkonjak, M. Gardner, V. Callaghan, P. Mattila, C. Guetl, V. M. Petrovic, K. Jovanovic, 'Virtual laboratories for education in science, technology, and engineering: A review,' Computers & Education, vol. 95, pp. 309-327, 2016. DOI: 10.1016/j.compedu.2016.02.002

[28] M. Irimia-Vladu, 'Green electronics: biodegradable and biocompatible materials and devices for sustainable future,' Chemical Society Reviews, vol. 43, no. 2, pp. 588-610, 2014. DOI: 10.1039/c3cs60235d

[29] G. J. Snyder, E. S. Toberer, 'Complex thermoelectric materials,' Nature Materials, vol. 7, no. 2, pp. 105-114, 2008. DOI: 10.1038/nmat2090

[30] Y. LeCun, Y. Bengio, G. Hinton, 'Deep learning,' Nature, vol. 521, no. 7553, pp. 436-444, 2015. DOI: 10.1038/nature14539

[31] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, et al., 'Loihi: A neuromorphic manycore processor with on-chip learning,' IEEE Micro, vol. 38, no. 1, pp. 82-99, 2018. DOI: 10.1109/mm.2018.112130359

[32] I. Afolabi, T. Taleb, K. Samdanis, A. Ksentini, H. Flinck, 'Network slicing and softwarization: A survey on principles, enabling technologies, and solutions,' IEEE Communications Surveys & Tutorials, vol. 20, no. 3, pp. 2429-2453, 2018. DOI: 10.1109/comst.2018.2815638

[33] X. Foukas, G. Patounas, A. Elmokashfi, M. K. Marina, 'Network slicing in 5G: Survey and challenges,' IEEE Communications Magazine, vol. 55, no. 5, pp. 94-100, 2017. DOI: 10.1109/mcom.2017.1600951

[34] P. Popovski, K. F. Trillingsgaard, O. Simeone, G. Durisi, '5G wireless network slicing for eMBB, URLLC, and mMTC: A communication-theoretic view,' IEEE Access, vol. 6, pp. 55765-55779, 2018. DOI: 10.1109/access.2018.2872781

[35] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, et al., 'A view of cloud computing,' Communications of the ACM, vol. 53, no. 4, pp. 50-58, 2010. DOI: 10.1145/1721654.1721672

[36] B. Varghese, R. Buyya, 'Next generation cloud computing: New trends and research directions,' Future Generation Computer Systems, vol. 79, pp. 849-861, 2018. DOI: 10.1016/j.future.2017.09.009

[37] J. Sweller, 'Cognitive load during problem solving: Effects on learning,' Cognitive Science, vol. 12, no. 2, pp. 257-285, 1988. DOI: 10.1016/0364-0213(88)90023-7

[38] F. Paas, A. Renkl, J. Sweller, 'Cognitive load theory and instructional design: Recent developments,' Educational Psychologist, vol. 38, no. 1, pp. 1-4, 2003. DOI: 10.1207/s15326985ep3801_1

[39] M. Tarafdar, Q. Tu, B. S. Ragu-Nathan, T. S. Ragu-Nathan, 'The impact of technostress on role stress and productivity,' Journal of Management Information Systems, vol. 24, no. 1, pp. 301-328, 2007. DOI: 10.2753/mis0742-1222240109

[40] T. S. Ragu-Nathan, M. Tarafdar, B. S. Ragu-Nathan, Q. Tu, 'The consequences of technostress for end users in organizations: Conceptual development and empirical validation,' Information Systems Research, vol. 19, no. 4, pp. 417-433, 2008. DOI: 10.1287/isre.1070.0165

[41] 'Under control: Audio/video conferencing systems feed surveillance capitalism with students data,' 2020 13th CMI Conference on Cybersecurity and Privacy (CMI), 2020. DOI: 10.1109/cmi51275.2020.9322736

[42] S. G. Hart, L. E. Staveland, 'Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research,' Advances in Psychology, vol. 52, pp. 139-183, 1988. DOI: 10.1016/s0166-4115(08)62386-9

[43] 'Surrogate gradient scaling for directly training spiking neural networks,' Applied Intelligence, 2023. DOI: 10.1007/s10489-023-04966-x