

An Offline-First Design Framework for Digital Learning in Low-Connectivity Environments

UCHECHI MARY-LINDA UNAMMA^{1*}, IFEANYICHUKWU JEFFREY OKWESA², UZOAMAKA IWUANYANWU³, DANIEL AJIGA⁴

¹Abia State University, Abia, Nigeria

²Tek Experts, Lagos State, Nigeria

³National Open University of Nigeria

⁴Alcorn State University, Mississippi

*Corresponding Author: Uchechi Mary-Linda Unamma

Abstract- Digital learning initiatives in the Global South are frequently designed on the tacit assumption of continuous, affordable, and high-bandwidth connectivity, an assumption that does not hold for the learners such initiatives most aspire to reach. When connectivity is intermittent, metered, and expensive, “connected-by-default” designs degrade into unusable ones, widening rather than narrowing educational inequality. This paper argues that offline operation should be treated not as a degraded fallback but as the primary, first-class design condition for learning technology in low-connectivity environments. Grounded in the literatures of information and communication technologies for development (ICT4D), human-computer interaction for development (HCI4D), mobile learning, and delay-tolerant networking, we propose TIDAL, a Tolerant, Inclusive, Device-first, Adaptive Learning framework. TIDAL couples five design tenets with a four-layer reference architecture (a local-first content and data layer, an opportunistic synchronization layer, an adaptive experience layer, and a pedagogical continuity layer), bound together by a cross-cutting concern for data integrity, security, and sustainability. The framework reframes connectivity as an intermittent resource to be exploited opportunistically rather than a precondition to be assumed. We illustrate TIDAL through a worked scenario of teacher-mediated secondary-school delivery in a rural setting and discuss the framework’s strengths, limitations, and boundary conditions. TIDAL is offered as a conceptual scaffold for designers, implementers, and evaluators, and as a research agenda for empirically validating offline-first learning design.

Keywords: offline-first; digital learning; low-connectivity; ICT4D; mobile learning; delay-tolerant networking; progressive web apps; design framework; digital inclusion

I. INTRODUCTION

The dominant grammar of contemporary educational technology is written in the language of the cloud: streaming video, real-time collaboration, continuous synchronization, and always-available content repositories. This grammar presupposes an infrastructure (dependable electricity, affordable data, and persistent broadband) that is unevenly distributed across the world. For a very large share of learners in low- and middle-income regions, connectivity is not a continuous utility but an intermittent, metered, and sometimes costly event [1], [2]. Where the design assumes the cloud and the reality delivers only occasional, fragile links, the gap between the two is not a minor inconvenience; it is the difference between a system that works and one that does not [3].

The character of the connectivity problem in these settings is easily misdescribed as one of simple scarcity, as though the only variable were the quantity of bandwidth available. In practice the problem is compound. Connectivity is not merely scarce but *variable*, present near a mast or a window and absent a few metres away, reliable at dawn and congested by evening, adequate for text and unequal to video. It is *metered*, so that every megabyte carries a price that the poorest users weigh against food, transport, and school fees, making the marginal cost of a failed or wasteful download a real deterrent to use. And it is *unpredictable*, so that a learner cannot plan around it as one plans around a scheduled power cut. A design that assumes continuity fails not because bandwidth is universally low but because it is intermittent along

several dimensions at once, and each dimension defeats a different assumption buried in a connected-by-default architecture.

This paper starts from a simple observation with far-reaching design implications: for many intended beneficiaries of digital learning, *offline is the normal operating state and connectivity is the exception*. Yet most learning platforms invert this ordering, treating offline behaviour as an afterthought, a spinner, an error message, or an ungraceful failure. The consequence is a form of exclusion that is easy to overlook because it is invisible to designers working in well-connected settings. A platform that technically “supports” a region but is unusable whenever the network falters does not extend access; it advertises it while withholding it [4].

This failure is peculiarly hard for well-resourced designers to perceive, and the difficulty is structural rather than a matter of carelessness. Development, testing, and demonstration typically occur on fast, unmetered office networks, where the offline path is exercised rarely if ever and its defects never surface. Usage analytics, gathered only when devices are online, systematically under-report the experience of users who are most often offline, rendering the excluded invisible in precisely the data meant to represent them. The learner who abandons a platform after it fails at low tide leaves no trace in a dashboard that counts only completed sessions. The result is a self-confirming blindness: the metrics that reach decision-makers describe the connected experience of the connected, while the degraded experience that drives exclusion goes unmeasured and therefore unaddressed. Naming offline as the design default is, among other things, a way of forcing this hidden experience into view.

We contend that closing this gap requires more than incremental engineering fixes such as caching a few assets or queuing a form submission. It requires a *design stance* in which offline capability is the default and connectivity is an opportunistic enhancement. This stance, commonly termed “offline-first” in contemporary web engineering, and closely related to the older networking tradition of delay- and disruption-tolerant systems [5], [6], has been articulated primarily as a set of technical patterns. What is missing, and what this paper supplies, is a

design framework that integrates the technical patterns with the pedagogical and human-centred requirements specific to *learning in resource-constrained* settings.

The distinction between an engineering fix and a design stance is central to the argument, and it is worth stating carefully. Caching an asset or queuing a submission are local remedies, each patching a particular point of failure while leaving the system’s underlying assumption, that the network is normally present, intact. A design stance, by contrast, revises that assumption at the root, so that every subsequent decision, from data storage to interface language to curriculum sequencing, is taken on the premise that the network is normally absent. The difference is not one of degree but of orientation: the same features that appear as reluctant concessions in a connected-by-default system become the natural, load-bearing structure of an offline-first one. A framework is needed precisely because this reorientation touches every layer of a learning system at once, and piecemeal fixes cannot deliver the coherence that a genuinely offline-tolerant experience requires.

The contribution of this paper is threefold. First, we synthesize four literatures that rarely speak to one another (ICT4D/HCI4D, mobile learning, delay-tolerant networking, and offline-first web engineering) into a coherent problem statement for offline-first learning design. Second, we propose TIDAL (*Tolerant, Inclusive, Device-first, Adaptive Learning*), a named conceptual framework comprising five design tenets and a four-layer reference architecture, with data integrity and resilience as a cross-cutting concern. Third, we demonstrate the framework’s explanatory and generative value through a worked scenario and a structured discussion of its strengths, limitations, and boundary conditions.

These contributions are pitched at three audiences whose needs the framework tries to serve simultaneously. For designers, TIDAL offers a stance and a checklist of load-bearing decisions that a connectivity-by-default habit would defer or miss. For implementers, it offers a layered vocabulary in which to allocate responsibilities and reason about trade-offs across the technical, interactional, and pedagogical dimensions of a system. For researchers and evaluators, it offers a set of testable propositions and an audit lens against which existing platforms can be

assessed. That a single scaffold can address these audiences at once is a consequence of its integrative ambition: because the problem of offline-first learning is inherently multidisciplinary, a framework that names its dimensions coherently is useful to everyone who must act within one of them.

There is, finally, an equity argument that gives the whole enterprise its urgency. The learners most likely to be failed by connectivity-by-default design are, almost by definition, those already least served by existing educational provision: the rural, the poor, the peripheral. A design default that quietly excludes them does not leave the distribution of educational opportunity unchanged; it reproduces and can even sharpen existing inequality under the banner of expanding access. To treat offline as the primary condition is therefore not merely a technical preference but a commitment to designing for the excluded rather than around them. The stakes of the default, in other words, are distributive: whichever operating state a system treats as normal, it implicitly treats the users who inhabit the other state as marginal, and in low-connectivity settings the marginalised state is the one the most disadvantaged learners occupy most of the time.

The paper is conceptual rather than empirical. Its aim is to offer a defensible, literature-grounded scaffold that designers can adopt, implementers can operationalize, and researchers can test. As with any framework, its ultimate value will be established through use and empirical scrutiny; we return to this point in the sections on discussion and future work.

A word is in order about the epistemic status of a conceptual contribution of this kind, since it conditions how the paper should be read and judged. TIDAL is not offered as an empirical finding but as a design theory: a structured, defeasible proposal about how the elements of an offline-first learning system should be organised and why. Its warrant comes from the convergence of four independent literatures on the inadequacy of connectivity-by-default and from the internal coherence with which it integrates their prescriptions, not from measured outcomes it does not yet possess. This means the framework is to be assessed, in the first instance, by whether it names the right problems and orders them defensibly, and only later, through the empirical programme sketched in

Section 6, by whether systems built on it demonstrably outperform their connectivity-dependent alternatives. Holding these two standards distinct guards against both premature dismissal and premature confidence.

Contributions on semiconductor, RF/MEMS and microelectronic systems engineering and biomedical imaging inform the framing adopted here [7]. Cognate work addressing enterprise IT systems, cloud migration, and service management provides complementary grounding [8].

II. BACKGROUND AND RELATED WORK

2.1 The design–reality gap in ICT4D

A recurring finding across two decades of ICT4D scholarship is that technology projects fail not because the technology is defective in the laboratory but because the assumptions embedded in the design diverge from the realities of the deployment context. Heeks [3] frames this as a shift from “ICT4D 1.0,” which imported designs and expectations wholesale from industrialized settings, toward “ICT4D 2.0,” which foregrounds context, appropriateness, and the active role of the user. Toyama [9], [10] sharpens the point with his “amplification” thesis: technology amplifies existing human and institutional capacity rather than substituting for it, so a poorly matched intervention amplifies dysfunction as readily as it amplifies benefit. Connectivity assumptions are a paradigmatic instance of the design–reality gap: a system architected for continuous bandwidth, deployed where bandwidth is intermittent, encodes a mismatch that no amount of user enthusiasm can overcome.

The design–reality gap is best understood not as a single mismatch but as a family of them, each corresponding to a dimension along which an imported design encodes an assumption the deployment context violates. Assumptions about infrastructure (continuous power, affordable bandwidth) are the most visible, but they are joined by assumptions about the user’s literacy and prior exposure to digital conventions, about the availability of technical support and maintenance, about the stability of the institutional environment, and about the very purposes the technology is meant to serve. Connectivity is the dimension this paper foregrounds

because it is at once pervasive, consequential, and unusually tractable: unlike institutional dysfunction or entrenched inequality, the assumption of continuous connectivity can be relaxed through design, and doing so removes one of the largest and most avoidable sources of mismatch. To say that a gap is designable-away is not to say it is trivial; it is to identify a point at which careful design can actually move the outcome, which is not true of every dimension of the reality the amplification thesis describes.

2.2 HCI for development and the marginal user

Human-computer interaction for development (HCI4D) extends this critique to the level of the interface and the interaction [11]. Its central insight is that the “average” user assumed by mainstream design (literate, experienced with graphical conventions, operating a current-generation device on a fast network) is not representative of many users in developing regions. Medhi, Sagar, and Toyama [12] demonstrate that text-heavy interfaces systematically exclude low-literacy users and that carefully designed text-free or voice-augmented interfaces can dramatically improve usability. The lesson generalizes: designing for the *marginal* device and the *marginal* user (the oldest handset, the smallest screen, the weakest signal, the least textual fluency) is not a niche accommodation but a precondition for genuine inclusion.

Designing for the marginal user has a counterintuitive property that recommends it beyond the population it directly serves: accommodations built for the most constrained case frequently improve the experience of everyone else. An interface legible to a low-literacy user through iconography and audio is also faster to parse for a fluent one; a system that functions on the weakest signal is more responsive on a strong one; a footprint small enough for the oldest handset leaves headroom on newer devices. This is the same logic that has made accessibility features in mainstream computing benefit users who were never their intended target. The implication for offline-first learning is that designing for the margin is not an act of charitable compromise that sacrifices quality for reach, but a discipline that tends to produce more robust systems across the board. The marginal user, in this sense, is not a special case to be handled after the general one but the design’s proper reference point.

2.3 Mobile learning in constrained contexts

Mobile learning research provides the pedagogical counterpart to these infrastructural and interactional concerns. Sharples, Taylor, and Vavoula [13] theorize learning as a mobile, conversational, and contextual process rather than a fixed transaction bound to a classroom or a desktop. Traxler [14] cautions that mobile learning cannot be understood merely as e-learning delivered on smaller screens; it entails distinct assumptions about place, time, and the fragmentation of attention and connectivity. In developing regions specifically, the mobile phone is frequently the *only* computing device a learner will ever own or access [2], and the “mobile-only” internet user faces constraints (data cost, battery, storage, intermittent coverage) that shape what is pedagogically feasible [1]. Anderson’s [15] synthesis of online-learning theory and practice reminds us that interaction (among learners, with teachers, and with content) is the pedagogical core that any delivery architecture must protect; UNESCO’s policy guidance similarly positions mobile technologies as a means to educational ends rather than an end in themselves [16]. Empirical projects such as MILLEE’s after-school literacy games in rural India [17] and Mitra’s [18] “minimally invasive education” experiments show both the promise of learner-driven technology use and the fragility of interventions that assume infrastructure the setting does not reliably provide.

What the mobile-learning literature contributes to the present argument is a reminder that the substrate and the pedagogy cannot be designed independently of one another. The properties of the delivery channel (its intermittency, its metering, the fragmentation of attention it imposes) are not neutral with respect to what can be taught and how. A pedagogy premised on continuous access to a live tutor or an open web of resources presupposes a channel that low-connectivity settings do not supply, and no amount of clever engineering at the delivery layer can fully reconstitute an interaction the channel cannot carry in real time. Conversely, a pedagogy designed from the outset around self-sufficient units, asynchronous feedback, and human mediation can thrive on an intermittent channel because its assumptions match the channel’s affordances. The lesson is that offline-first is not solely, or even primarily, a technical matter; it is a

pedagogical commitment to designs whose learning value does not depend on interactions the network cannot reliably deliver. This is why the framework proposed below refuses to treat the delivery architecture and the learning design as separable concerns.

2.4 Delay-tolerant networking and offline-first engineering

On the technical side, two traditions supply the mechanisms that make offline-first feasible. Delay- and disruption-tolerant networking (DTN) was conceived precisely for “challenged internets” where continuous end-to-end paths cannot be assumed; it replaces the expectation of an always-on connection with a *store-carry-forward* model in which data is held locally and relayed opportunistically when a link appears [5], [6]. More recently, offline-first web engineering, exemplified by progressive web apps (PWAs) built on service workers and local storage, has brought equivalent patterns to the mainstream browser, enabling applications that load and function without a network and synchronize when connectivity returns [19]. Brewer et al. [20] had already argued that technology for developing regions demands purpose-built architectures rather than hand-me-downs, and offline-first is the architectural expression of that argument for the learning domain. Selwyn [21], finally, provides a needed corrective: technology alone does not resolve educational inequities, and framework designers must remain honest about the social, institutional, and economic conditions that no architecture can engineer away.

It is worth drawing out the conceptual kinship between the two technical traditions this section names, because the framework depends on it. Delay-tolerant networking arose in contexts (deep-space communication, sparse sensor networks, disaster zones) where the assumption of a continuous end-to-end path was simply untenable, and its central move was to relocate reliability from the network to the endpoints: data is held, carried, and forwarded whenever a link happens to appear, and the absence of a link is treated as an ordinary operating condition rather than a fault. Offline-first web engineering makes the same move in a different idiom, relocating application state and logic to the client so that the application is complete on the device and treats the

server as an intermittent enhancement to be reconciled with opportunistically. Both traditions, in other words, invert the same assumption: that the network is normally present and its absence exceptional. The framework proposed here adopts that inversion as its foundation and asks what follows when it is applied not to packets or to generic applications but to learning, with all the pedagogical and human obligations that learning entails.

2.5 The gap this paper addresses

Each of these literatures is mature in isolation, but they are rarely integrated. DTN and offline-first engineering describe *how* to move and store data under intermittent connectivity but say little about pedagogy or the marginal learner. Mobile-learning theory describes *what* good learning looks like but generally assumes the delivery substrate as given. ICT4D/HCI4D name the *contextual* pitfalls but stop short of a delivery-architecture blueprint. What is missing is a framework that binds the technical patterns of offline-first operation to the pedagogical and human-centred requirements of learning in low-connectivity settings, and that makes data integrity and resilience, an operational-risk concern often ignored until it fails, an explicit design dimension. TIDAL is our attempt to fill that gap.

The integration this gap demands is not additive but architectural. It is not enough to place a delay-tolerant transport beneath a mobile-learning pedagogy beneath an HCI4D interface and declare the problem solved, because the requirements of each layer constrain the others: the pedagogy’s tolerance for asynchrony determines what the transport must guarantee, the transport’s reconciliation behaviour determines what the interface must honestly communicate, and the marginal device’s limits bound what content the pedagogy can assume. A framework worth the name must therefore specify not only the layers but the relations among them, and must elevate the concerns (data integrity, security, sustainability) that fall between the layers and are, for that reason, most often neglected. It is this insistence on the connective structure, rather than any single component, that distinguishes the contribution attempted here from the mature but disjoint literatures on which it draws.

Contributions on foundational information-society, ICT-for-development, digital-learning and forecasting literature inform the framing adopted here [22]–[55]. Cognate work addressing public-sector accounting, audit, and financial reporting provides complementary grounding [56].

III. THE TIDAL FRAMEWORK

3.1 Overview and naming

We name the framework TIDAL (*Tolerant, Inclusive, Device-first, Adaptive Learning*) because the tidal metaphor captures its central intuition: connectivity in low-resource environments ebbs and flows like a tide, and a well-designed system should function fully at low tide while opportunistically drawing on the high tide when it arrives. TIDAL is not a product or a protocol; it is a conceptual scaffold comprising (a) five design tenets that express its stance, (b) a four-layer reference architecture that operationalizes the tenets, and (c) a cross-cutting concern (data integrity, security, and sustainability) that spans all layers.

The tidal metaphor is more than an aide-mémoire; it encodes a substantive claim about how such systems should relate to connectivity. A shoreline is not damaged by the retreat of the tide, because it is built to stand at low water and merely to be enriched at high; a jetty designed to function only at high tide is useless for much of the day and dangerous when the water drops. Learning systems for low-connectivity settings should be built like the shoreline, complete in themselves at low tide and opportunistically enriched when connectivity floods in, rather than like the jetty, whose usefulness is hostage to a condition that prevails only intermittently. Holding this image in view helps keep the framework's tenets from collapsing back into the connected-by-default habits they are meant to displace.

3.2 The five design tenets

Tenet 1, Offline is the default state, not a failure mode. The system must be fully usable with no network connection for its core learning functions. Connectivity is treated as an enhancement that unlocks secondary functions (fresh content, remote grading, teacher dashboards), never as a precondition for primary ones. This inverts the conventional ordering

in which offline is a degraded exception to an online norm [5], [19].

Tenet 2, Synchronize opportunistically; degrade gracefully. Following the store-carry-forward logic of delay-tolerant networking [6], the system holds all state locally and reconciles with remote services whenever a usable link appears, however brief, metered, or unreliable. When connectivity is absent or costly, the experience degrades along a designed continuum rather than breaking abruptly.

Tenet 3, Design for the marginal device and the marginal user. The baseline target is the least-capable device and the least-connected, least-textually-fluent user the deployment realistically serves [11], [12]. Storage frugality, low memory footprint, and low-literacy-friendly interaction are first-order requirements, not optional accessibility layers.

Tenet 4, Preserve pedagogical coherence across connectivity states. The learner's experience of a lesson, an assessment, or a learning pathway must remain coherent whether they are offline, intermittently connected, or fully connected. Interaction (with content, teachers, and peers) is protected as the pedagogical core [13], [15], with connectivity-dependent interactions deferred and reconciled rather than blocked.

Tenet 5, Treat data integrity and resilience as first-class. Because state lives primarily on the device and is reconciled opportunistically, the risks of data loss, corruption, conflicting edits, and unauthorized access are elevated. The framework therefore embeds integrity, confidentiality, conflict resolution, and recoverability as explicit design obligations rather than deployment-time afterthoughts.

Taken together, the five tenets are best read not as independent rules but as a single stance viewed from five angles. The first fixes the default state; the second governs the transition between states; the third fixes the reference user and device against which every decision is measured; the fourth insists that the learning experience, not merely the software, survive the transitions; and the fifth guards the state that all of this places at risk by relocating it to the device. Removing any one tenet undermines the others: a system that made offline the default but ignored the

marginal device would run out of storage on the handsets that matter; one that synchronised opportunistically but neglected data integrity would corrupt the very learner records its opportunism was meant to preserve; one that protected pedagogical coherence but not confidentiality would expose children's data on the shared and easily lost devices that predominate. The tenets are load-bearing precisely in their mutual dependence, which is why the framework presents them as a set rather than a menu.

3.3 The four-layer reference architecture

The tenets are realized through four architectural layers, described below and depicted conceptually in Figure 1.

Layer 1, Local-First Content and Data Layer. This is the foundation. All content a learner needs for core functions (lesson packages, media, assessment items) and all state a learner generates (responses, progress, notes) reside on the device in durable local storage. Content is delivered as self-contained, versioned "learning packs" that can be pre-loaded (by physical distribution, side-loading, or a one-time download at a connectivity point) and that function with zero subsequent network dependence. Storage is budgeted explicitly against the constraints of low-end devices, and packs are structured so that partial or incremental updates can be applied without re-downloading whole packages. This layer is the architectural embodiment of Tenets 1 and 3.

Layer 2, Opportunistic Synchronization Layer. Sitting above local storage, this layer manages the ebb and flow of connectivity. It maintains an outbound queue of learner-generated state to be uploaded (assessment results, analytics, messages) and an inbound queue of updates to be fetched (new packs, corrections, teacher feedback). It detects usable connectivity, prioritizes transfers under bandwidth and cost constraints (small, high-value payloads first), resumes interrupted transfers, and applies a defined conflict-resolution policy when the same state has diverged across device and server. Its design draws directly on delay-tolerant, store-carry-forward principles [5], [6] and realizes Tenet 2.

Two design problems within this layer deserve to be made explicit, because they are where opportunistic

synchronisation most often goes wrong. The first is prioritisation under scarcity: when a link appears that may vanish within seconds and costs money for every byte, the layer must decide what to move first, and the right policy sends the smallest, highest-value payloads (a compact record of completed assessments, an urgent correction) ahead of bulky, deferrable content. The second is conflict: because the same state may evolve independently on the device and on the server between synchronisations, the layer must embody an explicit, defensible rule for reconciling divergence rather than silently overwriting one version with another. Neither problem arises in a connected-by-default system, where state is continuously reconciled and scarcity never bites; both are unavoidable once offline is the default, which is why the framework insists that the synchronisation layer own them deliberately rather than leave them to the accidents of implementation.

Layer 3, Adaptive Experience Layer. This layer governs what the learner sees and does as connectivity varies. It renders a connectivity-aware interface that communicates system state honestly (what is available now, what is pending, what will happen when a link returns) without punishing the learner for being offline. It applies low-literacy and low-bandwidth interaction patterns (iconography, audio, and progressive disclosure) informed by HCI4D findings [12]. Crucially, it enforces graceful degradation: connectivity-dependent affordances are visibly deferred rather than presented as broken. This layer operationalizes Tenets 3 and 4 at the point of interaction.

The demand that the interface communicate system state honestly is more consequential than it may appear, because dishonest or opaque signalling is a principal way that intermittency erodes trust. A learner who submits an answer and sees only a spinner, or worse an error, cannot tell whether their work was lost, is waiting, or succeeded; repeated experiences of this kind teach the rational conclusion that the system is unreliable and not worth the scarce data it consumes. An adaptive experience layer earns trust instead by making the invisible legible: work done offline is shown as safely recorded and pending, deferred actions are shown as deliberately postponed rather than broken, and the return of connectivity is marked

by the arrival of what was waiting. Honest signalling costs little to implement and pays for itself in continued use, whereas its absence quietly converts a technically sound offline-first system into one that users abandon for reasons its designers never see.

Layer 4, Pedagogical Continuity Layer. The topmost layer ensures that the *learning design*, not merely the software, tolerates intermittency. It sequences curricula so that offline-completable activities form self-sufficient units; it defines which pedagogical interactions may be asynchronous (a teacher's grading of a quiz, a peer discussion thread) and reconciles them when connectivity permits; and it embeds the teacher or facilitator as a first-class mediator who can bridge connectivity gaps through physical presence [16], [17]. This layer expresses Tenet 4 at the level of curriculum and human mediation, ensuring the framework serves learning rather than merely serving data.

Positioning the teacher or facilitator as a first-class element of the architecture, rather than as an external user of it, is one of the framework's more consequential moves. In connectivity-dependent designs the human intermediary is at best an afterthought and at worst an obstacle to be automated away; in an offline-first learning system the facilitator is frequently the most reliable channel available, carrying data physically between the school and the point of connectivity, resolving the interactions that cannot be made asynchronous, and adapting the pacing of a lesson to the rhythms of a particular classroom. Designing for this role means giving facilitators tools suited to it and accounting for their time and incentives in the system's sustainability budget. A framework that treated the delivery of data as the whole problem would render this human contribution invisible; by making pedagogical continuity a distinct layer, TIDAL keeps it in view as a designed-for resource rather than an unmanaged dependency.

Cross-cutting concern, Integrity, Security, and Sustainability. Spanning all four layers is a governance dimension reflecting Tenet 5. It comprises: *data integrity* (versioning, checksums, and idempotent reconciliation so that opportunistic sync never silently corrupts state); *security and privacy* (encryption of locally stored learner data and

authenticated, minimal-footprint sync, recognizing that devices in shared or insecure settings are easily lost or accessed by others); *conflict governance* (explicit, auditable rules for resolving divergent edits); and *sustainability* (energy, cost, and maintenance budgets, acknowledging that a system that drains a battery, exhausts a data bundle, or cannot be maintained locally will not endure). This dimension imports the operational-resilience and technology-risk perspective too often absent from learning-technology design.

Two properties of this layered scheme deserve emphasis. First, the layers are ordered by dependence rather than by importance: the local-first data layer is foundational because everything above it presupposes that state resides durably on the device, but it is the pedagogical-continuity layer at the top that gives the whole edifice its purpose, since a system that moves data flawlessly while fragmenting the learning experience has failed at its actual task. Second, the cross-cutting concern is drawn as a column spanning all four layers precisely because integrity, security, and sustainability cannot be localised to any one of them: a corruption introduced during synchronisation manifests in the data layer, a privacy breach exploits the interface, and an unsustainable energy budget defeats the pedagogy by draining the devices learners depend on. Treating these as a governance dimension that spans the stack, rather than as features of any single layer, is what allows the framework to address the failure modes that fall between conventional architectural boundaries and are, for that reason, the ones most likely to be missed.

3.4 Figure 1 (described)

Figure 1. The TIDAL layered reference architecture. The figure is a vertical stack of four horizontal bands, read from bottom to top. The bottom band, *Local-First Content and Data Layer*, is drawn widest and anchored to a device icon, signifying that everything rests on device-resident storage; inside it sit labelled blocks for “versioned learning packs” and “learner state store.” Above it, the *Opportunistic Synchronization Layer* is depicted with two opposed arrows, an upward “outbound queue” and a downward “inbound queue”, passing through a dashed, intermittent boundary labelled “connectivity (intermittent),” visually conveying that data crosses

the network only when a link appears. The third band, the *Adaptive Experience Layer*, shows a device screen whose interface elements are tagged with three states (“available,” “pending sync,” and “deferred”) illustrating graceful degradation. The top band, the *Pedagogical Continuity Layer*, contains icons for learner, teacher/facilitator, and curriculum sequence, connected by dotted lines that continue across the connectivity boundary to denote asynchronous, reconciled interaction. Running vertically along the right-hand edge, spanning all four bands, is a shaded column labelled *Integrity · Security · Sustainability*, signifying the cross-cutting governance concern. A wave motif along the connectivity boundary reinforces the tidal metaphor: the system stands complete at low tide and enriches itself at high tide.

Contributions on cybersecurity, data governance, and IT-risk management inform the framing adopted here [57]. Cognate work addressing Salesforce platform engineering and CRM governance provides complementary grounding [58].

IV. ILLUSTRATIVE APPLICATION: TEACHER-MEDIATED SECONDARY LEARNING IN A RURAL SCHOOL

To show how TIDAL guides concrete design decisions, we walk through a worked, illustrative scenario. It is not an empirical case study but a structured thought experiment demonstrating the framework’s generativity.

Setting. Consider a rural secondary school serving learners whose households own basic Android smartphones but who have no home broadband and only sporadic, expensive cellular data. The school itself has grid-independent solar power and occasional mobile coverage that becomes reliable only near a particular hillside or when a teacher travels to the nearest town. The pedagogical goal is to deliver a mathematics and science curriculum with formative assessment and teacher feedback.

The setting is deliberately chosen to sit squarely within the framework’s boundary conditions: connectivity is intermittent but not wholly absent, a facilitator is present to bridge the gaps, and the curriculum is of a kind that can be meaningfully pre-packaged for offline use. These features are what make the scenario a fair

test of the framework’s generativity rather than a demonstration rigged for success; a setting with continuous cheap connectivity or with no facilitator at all would fall outside the conditions under which TIDAL claims to add value, and applying it there would illustrate its limits rather than its use. Grounding the walk-through in a concrete, boundary-appropriate case keeps the subsequent design decisions honest, since each must answer to the specific constraints of solar power, sporadic coverage, basic handsets, and a single teacher acting as the connective tissue of the whole arrangement.

Applying Layer 1 (Local-First Content and Data). The curriculum is authored as versioned learning packs, one per topic, each containing lesson text, compressed diagrams, short audio explanations, and a bank of practice and assessment items. Packs are sized to fit comfortably within the storage budget of a low-end handset and are initially distributed by side-loading from the teacher’s device or a school laptop, requiring no per-learner download. From the learner’s perspective, every lesson and quiz works immediately and completely with the phone in airplane mode. All responses and progress are written to durable local storage.

The design decisions this layer forces are already substantive and would be easy to defer under a connected-by-default assumption. Sizing packs against the storage budget of the least-capable handset in the cohort determines how finely the curriculum must be segmented and how aggressively media must be compressed; choosing side-loading over per-learner download reflects a judgement that the household’s data allowance is too scarce and too valuable to spend on bulk content delivery; and structuring packs so that a correction can be applied incrementally, without re-downloading a whole topic, anticipates the reality that fixes will arrive through the same narrow, intermittent channel as everything else. None of these decisions is exotic, but each is one that a design assuming continuous connectivity would neglect until it failed in the field.

Applying Layer 2 (Opportunistic Synchronization). Each learner device maintains an outbound queue of completed quiz results and an inbound queue for new packs and teacher feedback. When a learner’s phone happens upon coverage (or, more reliably, when the

teacher's device (acting as a local relay) reaches the town)) the synchronization layer transfers the smallest, highest-value payloads first: compact assessment results upstream, then any waiting feedback and pack updates downstream. Transfers resume where they left off if the link drops mid-transfer. Because a learner might attempt the same quiz on two occasions before syncing, the layer applies a defined conflict policy (for instance, retaining the most recent complete attempt while preserving the earlier one for audit).

Applying Layer 3 (Adaptive Experience). The interface never blocks a learner for being offline. Submitted quizzes are marked "pending sync" with a clear, non-alarming indicator; teacher feedback, once received, surfaces as "new." Connectivity-dependent actions (such as requesting fresh content) are shown as deferred, with an honest explanation that they will complete when a link is available. Navigation relies on icons and audio prompts so that lower-literacy learners are not gated behind dense text, following HCI4D guidance [12].

Applying Layer 4 (Pedagogical Continuity). The curriculum is sequenced so that each topic pack is a self-sufficient offline unit: a learner can complete a full cycle of instruction, practice, and self-check without any network. Interactions that genuinely require another person (the teacher's grading of open-response items, a moderated peer discussion) are designed as asynchronous exchanges that queue locally and reconcile opportunistically. The teacher is positioned as the primary connectivity bridge, physically carrying data to and from town, echoing the facilitator-centred model that has proven effective in constrained deployments [16], [17].

Applying the cross-cutting concern. Locally stored learner data is encrypted, mitigating the risk that a lost or shared phone exposes a child's records. Packs and results carry version identifiers and checksums so that reconciliation is idempotent and corruption is detectable. The energy and data budgets are explicit: sync payloads are minimized to conserve both battery and the household's scarce data allowance, and packs are authored so the school can update and maintain them locally without specialist support, a precondition for the system's longevity [21].

What the framework surfaced. Applying TIDAL forces decisions that a connectivity-by-default design would defer or miss: the storage budget of the marginal device, the conflict policy for divergent quiz attempts, the teacher's role as a physical relay, the encryption of at-rest learner data, and the sequencing of curriculum into offline-sufficient units. This is the framework's intended value, not to prescribe a single implementation, but to make the load-bearing questions explicit and to order them coherently.

It is instructive to notice how many of these surfaced questions are not technical at all but organisational and pedagogical. The teacher's role as a physical relay is a matter of workflow and incentive as much as of software; the conflict policy for divergent quiz attempts is ultimately a question about what counts as a fair record of a learner's work; the sequencing of curriculum into offline-sufficient units is a curriculum-design decision with pedagogical stakes. A framework that surfaced only the engineering questions would leave these equally load-bearing decisions to chance. TIDAL's claim to usefulness rests substantially on its refusal to stop at the technical boundary, drawing the human and pedagogical decisions into the same ordered scheme as the architectural ones so that they are decided deliberately rather than by default.

Contributions on financial analytics, audit, IT-risk and governance across the wider research portfolio inform the framing adopted here [59]–[72]. Cognate work addressing procurement and supply-chain delivery in energy and infrastructure settings provides complementary grounding [73], [74].

V. DISCUSSION

5.1 Strengths

The principal strength of TIDAL is that it *reframes the default*. By treating offline as the normal state, it aligns the design's assumptions with the deployment's reality, directly narrowing the design–reality gap that Heeks [3] identifies as the dominant cause of ICT4D failure. A second strength is *integration*: TIDAL binds technical patterns (store-carry-forward synchronization, local-first storage) to human-centred and pedagogical requirements (marginal-user interaction, pedagogical continuity) and to

operational-risk concerns (integrity, security, sustainability) within a single scaffold, whereas these normally live in separate literatures and separate teams. Third, the framework is *generative and diagnostic*: as the worked scenario shows, it surfaces the load-bearing design questions and orders them, functioning both as a design checklist and as an evaluation lens against which existing platforms can be audited. Fourth, it is *substrate-neutral*: although progressive web apps and native mobile apps are natural implementation vehicles [19], the tenets and layers do not mandate a specific technology and can be realized across changing platforms.

A fifth strength, implicit in the preceding four, is that the framework functions as a shared language across the disciplines a real deployment must coordinate. Engineers, instructional designers, and the field staff who mediate learning ordinarily speak past one another, each attending to the concerns its own literature foregrounds. By naming tenets and layers that cut across these divisions and by making the cross-cutting concern explicit, TIDAL gives a multidisciplinary team a common vocabulary in which to negotiate trade-offs that would otherwise be resolved unilaterally, and often invisibly, by whichever discipline happens to own a given decision. The value of a shared scaffold in coordinating heterogeneous expertise should not be underestimated in settings where the failure to align these perspectives is itself a recurrent cause of disappointment.

5.2 Limitations

TIDAL is a conceptual contribution and inherits the limitations of its genre. It has not been empirically validated; its claims of usefulness rest on argument and literature synthesis, not on measured outcomes. The framework does not itself resolve genuinely *interactive* real-time pedagogies (synchronous tutoring, live collaboration) which are intrinsically connectivity-dependent; TIDAL can only defer and reconcile such interactions asynchronously, and some pedagogical value is inevitably lost in that translation. Nor does the framework eliminate cost: local-first designs shift burdens onto device storage, onto content-authoring effort, and onto the human facilitators who bridge connectivity gaps, and these burdens are real. Finally, as Selwyn [21] and Toyama [10] both warn, no architecture can substitute for the

social and institutional capacity that ultimately determines whether learning improves; TIDAL is a necessary enabler at best, never a sufficient cause.

A further limitation concerns the burden the framework shifts rather than removes. Local-first operation does not make the costs of intermittency disappear; it relocates them, onto device storage, onto the labour of authoring content that must be self-sufficient offline, and above all onto the human facilitators who carry data across connectivity gaps and bridge the interactions the network cannot. Where these facilitators are scarce, overburdened, or unsupported, the framework's dependence on them becomes a point of fragility rather than strength. Honesty requires acknowledging that TIDAL improves the odds of a workable system by aligning its assumptions with reality, but it does so partly by making explicit demands on human and organisational resources that are themselves in short supply in the settings it targets. Whether those demands can be met sustainably is an empirical question the framework raises but cannot answer.

5.3 Boundary conditions

The framework is most valuable where connectivity is *intermittent but not wholly absent* and where a *facilitator or physical relay* can bridge the gaps, the conditions of the worked scenario. Its value diminishes at the extremes. Where connectivity is genuinely continuous and cheap, the additional complexity of opportunistic synchronization may not pay for itself. Where there is *no* connectivity and *no* physical means of periodic data movement, even opportunistic sync has nothing to exploit, and the system collapses to a purely local, non-updating tool, still useful, but stripped of its adaptive character. TIDAL also presupposes that content can be meaningfully pre-packaged for offline use; pedagogies that depend on open-ended, real-time information retrieval sit awkwardly within it. Naming these boundaries is itself a contribution, since frameworks applied outside their conditions of validity tend to disappoint.

These boundaries also imply a spectrum rather than a binary, along which a given deployment can be located. Toward one end lies the fully connected, cheaply served setting where offline-first machinery is

over-engineering; toward the other lies the wholly disconnected setting with no means of periodic data movement, where the framework collapses to a static local tool. TIDAL's distinctive value is concentrated in the broad middle of this spectrum, intermittent connectivity plus some human or physical means of bridging the gaps, which happens to describe a very large share of the settings that digital-learning initiatives actually aspire to reach. Locating a deployment on this spectrum is itself a useful first step, since it tells a designer how much of the framework's apparatus is warranted and how much would be cost without corresponding benefit.

5.4 Relationship to existing practice

TIDAL does not claim to invent the offline-first stance, which has independent roots in delay-tolerant networking [5], [6] and contemporary web engineering [19]. Its contribution is to *specialize and integrate* that stance for the learning domain and the low-resource context, and to add the pedagogical-continuity and governance dimensions that engineering-centred treatments omit. In this sense TIDAL is complementary to, rather than competitive with, the mobile-learning theories [13], [14] and the HCI4D design guidance [11], [12] on which it draws; it supplies the connective architecture those bodies of work presuppose.

Positioned this way, TIDAL is best seen as a specialisation that inherits the credibility of its parent traditions while accepting a narrower scope in exchange for greater specificity. The generality that makes delay-tolerant networking and offline-first engineering broadly applicable also leaves them silent on the questions a learning designer must answer, just as the richness of mobile-learning theory leaves it dependent on a delivery substrate it does not itself provide. By binding these bodies of work to one another around the concrete problem of learning under intermittency, the framework trades breadth for usefulness at the point of decision. It does not compete with its sources for territory; it occupies the intersection they leave unattended, and its contribution stands or falls on whether that intersection is real and consequential, which the persistent failure of connectivity-by-default learning systems in low-resource settings strongly suggests it is.

Used in reverse, the same intersection supplies a diagnostic. Confronted with an existing learning platform that underperforms in a low-connectivity deployment, an evaluator can walk the four layers and the cross-cutting concern as an audit checklist, asking at each whether the system treats offline as default, synchronises opportunistically, honours the marginal device, preserves pedagogical coherence, and governs the state it places at risk. Failures tend to cluster at the seams the framework was built to expose, an application that caches content but blocks assessment when offline, or one that synchronises reliably but corrupts records under conflict, or one that is technically offline-capable but abandons the learner to opaque error states. Turning the framework into an audit instrument in this way costs nothing beyond the discipline of applying it, and it converts a design theory into a practical tool for improving systems that already exist rather than only for building new ones.

Contributions on internal audit, risk governance, and financial-sector controls inform the framing adopted here [75]. Cognate work addressing the broader external academic literature on analytics, machine learning, security, and digital systems provides complementary grounding [76]–[118].

VI. FUTURE WORK

Three lines of future work follow directly. First and most important is *empirical validation*. The framework should be instantiated in a concrete deployment and evaluated against connectivity-by-default alternatives on measures of usability, learning continuity, data integrity, and cost, in settings that match its boundary conditions. Mixed-methods evaluation, combining usage analytics with the contextual, qualitative methods characteristic of HCI4D [11], would be most informative. Second is the development of *reference implementations and design patterns*: reusable pack formats, conflict-resolution policies, and connectivity-aware interface components that lower the cost of adopting TIDAL and make its tenets concrete for practitioners. Third is the *extension of the framework* to adjacent problems (offline-tolerant assessment integrity, privacy-preserving synchronization for minors' data, and the economics of facilitator-mediated relay networks) each of which the cross-cutting governance concern gestures toward but does not fully specify. A further line worth

pursuing is the articulation of *maturity levels*, allowing an organization to locate an existing platform on a spectrum from connectivity-dependent to fully offline-first and to plan incremental improvement.

Beneath these concrete lines of work lies a broader research question that the framework poses but does not settle: what is the actual distribution of value across TIDAL's layers, and how does it shift with context? It is plausible that in some settings the decisive contribution is the local-first data layer, which keeps the system usable at all, while the pedagogical-continuity layer adds comparatively little; in others the reverse may hold, with the human-mediated continuity of learning proving more consequential than any transport optimisation. Empirical study should therefore aim not merely to validate the framework as a whole against connectivity-by-default alternatives but to decompose its contribution, identifying which tenets and layers carry the most weight under which conditions. Such decomposition would convert the framework from a coherent proposal into a calibrated instrument, telling designers not only what to consider but where, in their particular circumstances, to concentrate scarce effort.

Contributions on financial analytics, planning, and enterprise decision systems inform the framing adopted here [119], [120]. Cognate work addressing the foundational and directly cited literature underpinning the present study provides complementary grounding [1]–[6], [9]–[21].

VII. CONCLUSION

Digital learning too often reaches the under-connected only in name, offering a technically supported experience that fails the moment the network does. This paper has argued that the remedy is not incremental patching but a change of default: offline operation must be treated as the first-class design condition, and connectivity as an opportunistic enhancement. We proposed TIDAL (a *Tolerant, Inclusive, Device-first, Adaptive Learning* framework) that couples five design tenets with a four-layer reference architecture and a cross-cutting concern for integrity, security, and sustainability, and we demonstrated through a worked scenario how the framework surfaces and orders the decisions that matter. TIDAL is grounded in, and integrates, the

ICT4D, HCI4D, mobile-learning, and delay-tolerant-networking literatures, and it is honest about its own limits: it is a conceptual scaffold awaiting empirical test, a necessary enabler rather than a sufficient cause of better learning. Offered in that spirit, it aims to give designers, implementers, and researchers a common language and a defensible starting point for building learning technology that works where learners actually are, at low tide as well as high.

The wager underlying the framework is finally a simple one: that aligning a system's assumptions with the conditions its users actually inhabit does more to extend genuine access than any amount of capability added on the assumption of conditions they do not. If that wager is correct, the path to reaching the under-connected runs not through ever more ambitious demands on infrastructure but through more disciplined humility about its absence, and the design default is the place where that humility must first be encoded. TIDAL is one attempt to encode it; the task now is to put it to the test.

REFERENCES

- [1] Gitau, S., Marsden, G., & Donner, J. (2010). After access: Challenges facing mobile-only Internet users in the developing world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)* (pp. 2603–2606). ACM. [Crossref](#)
- [2] Donner, J. (2015). *After access: Inclusion, development, and a more mobile Internet*. MIT Press.
<https://mitpress.mit.edu/9780262029926/after-access/>
- [3] Heeks, R. (2008). ICT4D 2.0: The next phase of applying ICT for international development. *Computer*, 41(6), 26–33. [Crossref](#)
- [4] Warschauer, M. (2003). *Technology and social inclusion: Rethinking the digital divide*. MIT Press.
<https://mitpress.mit.edu/9780262731737/technology-and-social-inclusion/>
- [5] Fall, K. (2003). A delay-tolerant network architecture for challenged internets. In

- Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '03)* (pp. 27–34). ACM. [Crossref](#)
- [6] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., & Weiss, H. (2007). *Delay-tolerant networking architecture* (RFC 4838). Internet Engineering Task Force. [Crossref](#)
- [7] Ejofodomi, O. A., Gideon, E. N., Oladipo, G. O., & Oshomah, E. R. (2014). Automated detection of architectural detection in mammograms using template matching. *International Journal of Biomedical Science and Engineering*, 2(1), 1–6.
- [8] Ladapo, O. O., Dosunmu, A. A., Jooda, D., & Abolaji, T. O. (2018). Lessons learned from offline assessment of security-critical systems: The case of Microsoft Active Directory. *Iconic Research and Engineering Journals*, 2(6), 277–299. [Crossref](#)
- [9] Toyama, K. (2011). Technology as amplifier in international development. In *Proceedings of the 2011 iConference* (pp. 75–82). ACM. [Crossref](#)
- [10] Toyama, K. (2015). *Geek heresy: Rescuing social change from the cult of technology*. PublicAffairs.
- [11] Ho, M. R., Smyth, T. N., Kam, M., & Kotlarsky, A. (2009). Human-computer interaction for development: The past, present, and future. *Information Technologies & International Development*, 5(4), 1–18. <https://itidjournal.org/index.php/itid/article/view/420>
- [12] Medhi, I., Sagar, A., & Toyama, K. (2007). Text-free user interfaces for illiterate and semiliterate users. *Information Technologies & International Development*, 4(1), 37–50. [Crossref](#)
- [13] Sharples, M., Taylor, J., & Vavoula, G. (2005). Towards a theory of mobile learning. In *Proceedings of mLearn 2005: 4th World Conference on mLearning*. Cape Town, South Africa. https://www.researchgate.net/publication/228346088_Towards_a_theory_of_mobile_learning
- [14] Traxler, J. (2007). Defining, discussing and evaluating mobile learning. *International Review of Research in Open and Distributed Learning*, 8(2). [Crossref](#)
- [15] Anderson, T. (Ed.). (2008). *The theory and practice of online learning* (2nd ed.). Athabasca University Press. <https://www.aupress.ca/books/120146-the-theory-and-practice-of-online-learning/>
- [16] West, M., & Vosloo, S. (2013). *UNESCO policy guidelines for mobile learning*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000219641>
- [17] Kam, M., Kumar, A., Jain, S., Mathur, A., & Canny, J. (2009). Improving literacy in rural India: Cellphone games in an after-school program. In *Proceedings of the 3rd International Conference on Information and Communication Technologies and Development (ICTD 2009)* (pp. 139–149). IEEE. [Crossref](#)
- [18] Mitra, S. (2003). Minimally invasive education: A progress report on the “hole-in-the-wall” experiments. *British Journal of Educational Technology*, 34(3), 367–371. [Crossref](#)
- [19] Biørn-Hansen, A., Majchrzak, T. A., & Grønli, T.-M. (2017). Progressive web apps: The possible web-native unifier for mobile development. In *Proceedings of the 13th International Conference on Web Information Systems and Technologies (WEBIST 2017)* (pp. 344–351). SciTePress. [Crossref](#)
- [20] Brewer, E., Demmer, M., Du, B., Ho, M., Kam, M., Nedeveschi, S., Pal, J., Patra, R., Surana, S., & Fall, K. (2005). The case for technology in developing regions. *Computer*, 38(6), 25–38. [Crossref](#)
- [21] Selwyn, N. (2016). *Education and technology: Key issues and debates* (2nd ed.). Bloomsbury Academic.
- [22] Compaine, B. M. (Ed.). (2001). *The digital divide: Facing a crisis or creating a myth?* MIT Press.

- [23] Mossberger, K., Tolbert, C. J., & Stansbury, M. (2003). *Virtual inequality: Beyond the digital divide*. Georgetown University Press.
- [24] Attewell, P. (2001). The first and second digital divides. *Sociology of Education*, 74(3), 252–259. [Crossref](#)
- [25] Helsper, E. J. (2012). A corresponding fields model for the links between social and digital exclusion. *Communication Theory*, 22(4), 403–426. [Crossref](#)
- [26] van Deursen, A. J. A. M., & Helsper, E. J. (2015). The third-level digital divide: Who benefits most from being online? *Communication and Information Technologies Annual*, 10, 29–52. [Crossref](#)
- [27] Scheerder, A., van Deursen, A., & van Dijk, J. (2017). Determinants of Internet skills, uses and outcomes: A systematic review of the second- and third-level digital divide. *Telematics and Informatics*, 34(8), 1607–1624. [Crossref](#)
- [28] Heeks, R. (2010). Do information and communication technologies (ICTs) contribute to development? *Journal of International Development*, 22(5), 625–640. [Crossref](#)
- [29] Walsham, G. (2017). ICT4D research: Reflections on history and future agenda. *Information Technology for Development*, 23(1), 18–41. [Crossref](#)
- [30] Gunkel, D. J. (2003). Second thoughts: Toward a critique of the digital divide. *New Media & Society*, 5(4), 499–522. [Crossref](#)
- [31] DiMaggio, P., Hargittai, E., Neuman, W. R., & Robinson, J. P. (2001). Social implications of the Internet. *Annual Review of Sociology*, 27, 307–336. [Crossref](#)
- [32] Wei, K.-K., Teo, H.-H., Chan, H. C., & Tan, B. C. Y. (2011). Conceptualizing and testing a social cognitive model of the digital divide. *Information Systems Research*, 22(1), 170–187. [Crossref](#)
- [33] Hilbert, M. (2011). The end justifies the definition: The manifold outlooks on the digital divide and their practical usefulness for policy-making. *Telecommunications Policy*, 35(8), 715–736. [Crossref](#)
- [34] Selwyn, N. (2016). *Is technology good for education?* Polity Press.
- [35] Sharples, M., Taylor, J., & Vavoula, G. (2010). A theory of learning for the mobile age. In *Medienbildung in neuen Kulturräumen* (pp. 87–99). VS Verlag. [Crossref](#)
- [36] Traxler, J. (2007). Defining, discussing and evaluating mobile learning: The moving finger writes and having writ... *The International Review of Research in Open and Distributed Learning*, 8(2), 1–12. [Crossref](#)
- [37] Kukulska-Hulme, A. (2009). Will mobile learning change language learning? *ReCALL*, 21(2), 157–165. [Crossref](#)
- [38] Donner, J. (2008). Research approaches to mobile use in the developing world: A review of the literature. *The Information Society*, 24(3), 140–159. [Crossref](#)
- [39] Brown, J. S., & Adler, R. P. (2008). Minds on fire: Open education, the long tail, and Learning 2.0. *EDUCAUSE Review*, 43(1), 16–32.
- [40] Kozma, R. B. (2003). Technology and classroom practices: An international study. *Journal of Research on Technology in Education*, 36(1), 1–14. [Crossref](#)
- [41] Pittinsky, M. S. (2003). *The wired tower: Perspectives on the impact of the Internet on higher education*. Financial Times Prentice Hall.
- [42] Box, G. E. P., & Jenkins, G. M. (1976). *Time series analysis: Forecasting and control* (Rev. ed.). Holden-Day.
- [43] Armstrong, J. S. (Ed.). (2001). *Principles of forecasting: A handbook for researchers and practitioners*. Kluwer Academic. [Crossref](#)

- [44] De Gooijer, J. G., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International Journal of Forecasting*, 22(3), 443–473. [Crossref](#)
- [45] Gardner, E. S. (2006). Exponential smoothing: The state of the art: Part II. *International Journal of Forecasting*, 22(4), 637–666. [Crossref](#)
- [46] Holt, C. C. (2004). Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, 20(1), 5–10. [Crossref](#)
- [47] Fildes, R., Nikolopoulos, K., Crone, S. F., & Syntetos, A. A. (2008). Forecasting and operational research: A review. *Journal of the Operational Research Society*, 59(9), 1150–1172. [Crossref](#)
- [48] Syntetos, A. A., Boylan, J. E., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, 56(5), 495–503. [Crossref](#)
- [49] Ahmed, N. K., Atiya, A. F., Gayar, N. E., & El-Shishiny, H. (2010). An empirical comparison of machine learning models for time series forecasting. *Econometric Reviews*, 29(5–6), 594–621. [Crossref](#)
- [50] Crone, S. F., Hibon, M., & Nikolopoulos, K. (2011). Advances in forecasting with neural networks? Empirical evidence from the NN3 competition. *International Journal of Forecasting*, 27(3), 635–660. [Crossref](#)
- [51] Makridakis, S., Wheelwright, S. C., & Hyndman, R. J. (1998). *Forecasting: Methods and applications* (3rd ed.). Wiley.
- [52] Fildes, R., & Goodwin, P. (2007). Against your better judgment? How organizations can improve their use of management judgment in forecasting. *Interfaces*, 37(6), 570–576. [Crossref](#)
- [53] Chatfield, C. (2000). *Time-series forecasting*. Chapman & Hall/CRC.
- [54] Hyndman, R. J., Koehler, A. B., Snyder, R. D., & Grose, S. (2002). A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of Forecasting*, 18(3), 439–454. [Crossref](#)
- [55] Norris, P. (2001). *Digital divide: Civic engagement, information poverty, and the Internet worldwide*. Cambridge University Press. [Crossref](#)
- [56] Dogbatsey, E. A., & Ebhojie, O. (2018). Budget compliance and statutory reporting in Sub-Saharan Africa: A systematic review of frameworks, gaps, and reform pathways. Zenodo. [Crossref](#)
- [57] Mbonu, I. S., Aliliele, C., Iwuanyanwu, U., & Oluoha, O. M. (2018). A conceptual framework for legal and ethical risk modeling in enterprise data protection governance systems. *Iconic Research and Engineering Journals*, 2(2), 207–226. [Crossref](#)
- [58] Badmus, O., Dosunmu, A. A., & Ozowara, D. E. (2018). A systematic review of CI/CD pipeline strategies in Salesforce DevOps: Tools, practices, and deployment outcomes. *Iconic Research and Engineering Journals*, 2(6).
- [59] Akeju, B., Edivri, J., Ogbole, J. I., Okoruwa, P. O., Fadayomi, O., & Abolaji, T. O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. *Iconic Research and Engineering Journals*, 1(9), 476–492. [Crossref](#)
- [60] Aminu-Ibrahim, A. Y., Ogbete, J. C., & Ambali, K. B. (2018). Developing sustainable diagnostic laboratory infrastructure models for emerging and resource constrained health systems. *Iconic Research and Engineering Journals*, 1(8), 118–132. [Crossref](#)
- [61] Ogbete, J. C., Aminu-Ibrahim, A. Y., & Ambali, K. B. (2018). Optimizing laboratory spatial planning strategies to improve diagnostic accuracy, safety, and clinical throughput. *Iconic Research and Engineering Journals*, 2(1), 87–113. [Crossref](#)
- [62] Arumosoye, O. M., & Obriki, O. D. (2018). Development of an integrated heat stress risk conceptual model for industrial operations in

- extreme environments. *Iconic Research and Engineering Journals*, 1(12), 141-160. [Crossref](#)
- [63] Obriki, O. D., & Arumosoye, O. M. (2018). Conceptual modeling of data-driven occupational safety risk control in large-scale energy infrastructure projects. *Iconic Research and Engineering Journals*, 1(7), 169-189. [Crossref](#)
- [64] Dagodzo, D. (2018b). A conceptual framework for UAV integration into national power grid inspection programs. *Iconic Research and Engineering Journals*, 2(5), 391-412. [Crossref](#)
- [65] Dagodzo, D. (2018c). A review of UAV applications in electrical transmission line inspection: Methods, technologies, and challenges. *Iconic Research and Engineering Journals*, 2(6), 234-254. [Crossref](#)
- [66] Sunday, E. A., & Omoegun, G. O. (2018). Integrating solar power solutions in small-scale manufacturing industries in Nigeria. *International Journal of Scientific Research in Science, Engineering and Technology*, 4(8), 832-853.
- [67] Bobga, M. A., Boakye, K., Ogbona, C. S., & Yeboah, T. J. (2018). Pedagogical strategies for teaching students with learning difficulties in resource-constrained schools. *Iconic Research and Engineering Journals*, 2(4), 173-194. [Crossref](#)
- [68] Amayo, E. B. (2017). Multi algorithm of loss objective function of a single phase induction motor. *International Journal of Development Research*, 7(5), 12805-12810.
- [69] Amayo, E. B., & Popoola, T. T. (2017a). Multi algorithm of weight objective function of a single phase induction motor. *International Journal of Trend in Research and Development*, 4(2), 184-188.
- [70] Amayo, E. B., & Popoola, T. T. (2017b). Scientific study of telecommunication deployment in achieving quality network. *International Journal of Trend in Research and Development*, 4(5), 311-314.
- [71] Amayo, E. B., Ubeku, E., & Oshevire, P. (2015). Multi algorithm of a single objective function of a single phase induction motor. *Journal of Multidisciplinary Engineering Science and Technology*, 2(12), 3400-3403.
- [72] Oshevire, P., Eyenubo, O. J., & Amayo, B. (2017). Voltage control in the presence of distributed generation. *ATBU Journal of Science, Technology and Education*, 5(2), 165-173.
- [73] Okonkwo, C. S., Ogunwole, O., & Okeke, O. T. (2018b). Framework for strategic procurement optimization in oil and gas operations. *Iconic Research and Engineering Journals*, 1(7), 153-168. [Crossref](#)
- [74] Okonkwo, C. S., Ogunwole, O., & Okeke, O. T. (2018c). Model for inventory availability and plant uptime improvement in energy facilities. *Iconic Research and Engineering Journals*, 2(4), 160-172. [Crossref](#)
- [75] Akomolafe, O., & Agu, M. U. (2018). A conceptual model for enhancing internal audit quality through technology-enabled risk assessment frameworks. *Iconic Research and Engineering Journals*, 1(9), 458-475.
- [76] Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7), 1483-1510. [Crossref](#)
- [77] Susto, G. A., Schirru, A., Pampuri, S., McLoone, S., & Beghi, A. (2015). Machine learning for predictive maintenance: A multiple classifier approach. *IEEE Transactions on Industrial Informatics*, 11(3), 812-820. [Crossref](#)
- [78] Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. *Journal of Manufacturing Systems*, 48, 144-156. [Crossref](#)
- [79] Stouffer, K., Lightman, S., Pillitteri, V., Abrams, M., & Hahn, A. (2015). Guide to Industrial Control Systems (ICS) Security. NIST SP 800-82.

- [80] Dragos, & Inc. (2017). TRISIS Malware: Analysis of Safety System Targeted Attack. Dragos Intelligence.
- [81] Papernot, N., McDaniel, P., Sinha, A., & Wellman, M. P. (2018). SoK: Security and privacy in machine learning. in Proc. IEEE EuroS&P, 399-414.
- [82] Sproul, W. D., Christie, D. J., & Carter, D. C. (2005). Control of reactive sputtering processes. *Thin Solid Films*, 491(1), 1–17. [Crossref](#)
- [83] Strijckmans, K., Schelfhout, R., & Depla, D. (2018). Tutorial: Hysteresis during the reactive magnetron sputtering process. *Journal of Applied Physics*, 124(24), 241101. [Crossref](#)
- [84] Musil, J., Baroch, P., Vlček, J., Nam, K. H., & Han, J. G. (2005). Reactive magnetron sputtering of thin films: present status and trends. *Thin Solid Films*, 475(1), 208–218. [Crossref](#)
- [85] Sarakinos, K., Alami, J., & Konstantinidis, S. (2010). High power pulsed magnetron sputtering: A review on scientific and engineering state of the art. *Surface and Coatings Technology*, 204(11), 1661–1684. [Crossref](#)
- [86] Anders, A. (2017). Tutorial: Reactive high power impulse magnetron sputtering (R-HiPIMS). *Journal of Applied Physics*, 121(17), 171101. [Crossref](#)
- [87] Anders, A. (2014). A review comparing cathodic arcs and high power impulse magnetron sputtering (HiPIMS). *Surface and Coatings Technology*, 257, 308–325. [Crossref](#)
- [88] Negri, E., Fumagalli, L., & Macchi, M. (2017). A Review of the Roles of Digital Twin in CPS-based Production Systems. *Procedia Manufacturing*, 11, 939–948. [Crossref](#)
- [89] Kritzing, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital Twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1016–. [Crossref](#)
- [90] Brunton, S. L., Proctor, J. L., & Kutz, J. N. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 113(15), 3932–3937. [Crossref](#)
- [91] Benner, P., Gugercin, S., & Willcox, K. (2015). A Survey of Projection-Based Model Reduction Methods for Parametric Dynamical Systems. *SIAM Review*, 57(4), 483–531. [Crossref](#)
- [92] Evensen, G. (2003). The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Ocean Dynamics*, 53(4), 343–367. [Crossref](#)
- [93] Mayne, D. Q. (2014). Model predictive control: Recent developments and future promise. *Automatica*, 50(12), 2967–2986. [Crossref](#)
- [94] Dagodzo, D. (2018a). A conceptual framework for UAV integration into national power grid inspection programs. *IRE Journals*, 2(5), 391–412. [Crossref](#)
- [95] Pflug, A., Siemers, M., Melzig, T., Schäfer, L., & Bräuer, G. (2014). Simulation of linear magnetron discharges in 2D and 3D. *Surface and Coatings Technology*, 260, 411–416. [Crossref](#)
- [96] Berg, S., Särhammar, E., & Nyberg, T. (2014). Upgrading the ‘Berg-model’ for reactive sputtering processes. *Thin Solid Films*, 565, 186–192. [Crossref](#)
- [97] Elebe, O. (2018). Conceptual model for insider threat classification and risk modeling in complex digital systems. and digital identity verification framework for.
- [98] Srinidhi, B., Yan, J., & Bhargava, H. K. (2015). Effect of information security investments on firm performance. *Decision Support Systems*, 74, 1–15.
- [99] McKone, K. E., Schroeder, R. G., & Cua, K. O. (2001). The impact of total productive maintenance practices on manufacturing performance. *Journal of Operations Management*, 19(1), 39–58. [Crossref](#)

- [100] Dal, B., Tugwell, P., & Greatbanks, R. (2000). Overall equipment effectiveness as a measure of operational improvement: A practical analysis. *International Journal of Operations & Production Management*, 20(12), 1488–1502. [Crossref](#)
- [101] Bamber, D., Sharp, C. J., & Hides, M. T. (1999). Factors affecting successful implementation of total productive maintenance: A UK manufacturing case study perspective. *Journal of Quality in Maintenance Engineering*, 5(3), 162–181. [Crossref](#)
- [102] McKone, K. E., Schroeder, R. G., & Cua, K. O. (1999). Total productive maintenance: A contextual view. *Journal of Operations Management*, 17(2), 123–144. [Crossref](#)
- [103] Tezel, A., Koskela, L., & Tzortzopoulos, P. (2016). Visual management in production management: A literature synthesis. *Journal of Manufacturing Technology Management*, 27(6), 766–799. [Crossref](#)
- [104] Bagavathiappan, S., Lahiri, B. B., Saravanan, T., Philip, J., & Jayakumar, T. (2013). Infrared thermography for condition monitoring: A review. *Infrared Physics & Technology*, 60, 35–55. [Crossref](#)
- [105] Lee, J., Wu, F., Zhao, W., Ghaffari, M., Liao, L., & Siegel, D. (2014b). Prognostics and health management design for rotary machinery systems: Reviews, methodology and applications. *Mechanical Systems and Signal Processing*, 42(1), 314–334. [Crossref](#)
- [106] Lei, Y., Li, N., Guo, L., Li, N., Yan, T., & Lin, J. (2018). Machinery health prognostics: A systematic review from data acquisition to RUL prediction. *Mechanical Systems and Signal Processing*, 104, 799–834. [Crossref](#)
- [107] Saxena, A., Goebel, K., Simon, D., & Eklund, N. (2008). Damage propagation modeling for aircraft engine run-to-failure simulation. in Proc. Int. Conf. Prognostics and Health Management (PHM), 1–9. [Crossref](#)
- [108] Lee, J., Kao, H.-A., & Yang, S. (2014a). Service innovation and smart analytics for Industry 4.0 and big data environment. *Procedia CIRP*, 16, 3–8. [Crossref](#)
- [109] Tao, F., Zhang, M., Liu, Y., & Nee, A. Y. C. (2018b). Digital twin driven prognostics and health management for complex equipment. *CIRP Annals*, 67(1), 169–172. [Crossref](#)
- [110] Selcuk, S. (2017). Predictive maintenance, its implementation and latest trends. *Proceedings of the Institution of Mechanical Engineers*, 231(9), 1670–1679. [Crossref](#)
- [111] Okonkwo, C. S., Ogunwale, O., & Okeke, O. T. (2018a). Model for inventory availability and plant uptime improvement in energy facilities. *IRE Journals*, 2(4), 160–172. [Crossref](#)
- [112] Peng, Y., Dong, M., & Zuo, M. J. (2010). Current status of machine prognostics in condition-based maintenance: A review. *The International Journal of Advanced Manufacturing Technology*, 50(1), 297–313. [Crossref](#)
- [113] Liu, H.-C., Liu, L., & Liu, N. (2013). Risk evaluation approaches in failure mode and effects analysis: A literature review. *Expert Systems with Applications*, 40(2), 828–838. [Crossref](#)
- [114] Antoni, J. (2007). Fast computation of the kurtogram for the detection of transient faults. *Mechanical Systems and Signal Processing*, 21(1), 108–124. [Crossref](#)
- [115] Lei, Y., Lin, J., Zuo, M. J., & He, Z. (2014). Condition monitoring and fault diagnosis of planetary gearboxes: A review. *Measurement*, 48, 292–305. [Crossref](#)
- [116] Lei, Y., Jia, F., Lin, J., Xing, S., & Ding, S. X. (2016). An intelligent fault diagnosis method using unsupervised feature learning towards mechanical big data. *IEEE Transactions on Industrial Electronics*, 63(5), 3137–3147. [Crossref](#)

- [117] Wen, L., Li, X., Gao, L., & Zhang, Y. (2018). A new convolutional neural network-based data-driven fault diagnosis method. *IEEE Transactions on Industrial Electronics*, 65(7), 5990–. [Crossref](#)
- [118] Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., & Sui, F. (2018a). Digital twin-driven product design, manufacturing and service with big data. *The International Journal of Advanced Manufacturing Technology*, 94(9), 3563–3576. [Crossref](#)
- [119] Lawal, O. A., & Oduleye, T. E. (2018a). A conceptual model for financial analytics driven enterprise value creation in technology firms. *Iconic Research and Engineering Journals*, 2(2).
- [120] Lawal, O. A., & Oduleye, T. E. (2018b). A review and conceptual framework for tax governance and cross-border compliance analytics. *Iconic Research and Engineering Journals*, 2(5).