WWW.IJRPR.COM

International Journal of Research Publication and Reviews

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

The Multifaceted Ecosystem of Mobile Telephony: An Analysis of Cell Phone Applications, Technological Misuses and Societal Misconceptions

[1]Chavan Samruddhi Santosh, [2]Chavan Sakshi Vinayak, [3]Jarhad Janhavi Kishor, [4]Swapnil Shivaji Sajane

[1,2,3]Student, Department of Mechanical Engineering, Ahilyanagar, India.
[4]Guide, Department of Mechanical Engineering, Ahilyanagar, India.
* *E-mail: chavhansamruddhi459@gmail.com*

ABSTRACT:

The contemporary cell phone has transcended its original role as a telecommunications device to become a ubiquitous computational platform capable of complex sensing, on-device inference and cross-modal retrieval. This paper presents a structured analysis of the mobile telephony ecosystem along two complementary axes: the expanding utility of consumer handset hardware, and the misuses—technical as well as societal—that constrain its safe and equitable deployment. On the utility axis, empirical literature is synthesised across four themes: low-cost scientific sensing that repurposes complementary metal-oxide-semiconductor (CMOS) camera sensors for cosmic-ray and ionising-radiation detection; cross-modal retrieval of Musical Instrument Digital Interface (MIDI) passages from photographs of printed sheet music; satellite-imagery-based forecasting of network adoption for digital-inclusion planning; and linguistically optimised keypad layouts for short message service (SMS) entry. On the misuse axis, the paper examines Application Programming Interface (API) misuse in the data-centric software libraries that underpin mobile analytics back ends, together with the persistent and scientifically unfounded public belief that handset radio-frequency emissions cause cancer, a claim incompatible with the energetics of the photoelectric effect. To connect these strands, a three-module framework—the Mobile Utility and Risk Assessment Framework (MURAF)—is proposed. It couples lightweight deterministic sensing, infrastructure and interface mapping, and large-language-model-assisted misuse detection with a public scientific-literacy component, and it is accompanied by a reproducible evaluation protocol with explicit target metrics. The analysis concludes that optimisation of the mobile ecosystem depends as much on software robustness and public scientific literacy as it does on further hardware innovation.

Keywords: Mobile Telephony; Sensor Repurposing; Cross-Modal Retrieval; API Misuse; Radio-Frequency Radiation; Digital Divide; Mobile Software Quality

1. Introduction

Mobile telephony is today the most widely distributed computing platform in history. Devices originally engineered for voice calls and Short Message Service (SMS) exchange now integrate multi-megapixel optical sensors, multi-core application processors, inertial and magnetic transducers, and heterogeneous connectivity modules spanning cellular, Wi-Fi, Bluetooth and satellite-navigation standards. This convergence of sensing, computation and communication has changed the analytical status of the handset: it is no longer merely a terminal through which information is consumed, but an instrument through which measurements are produced. The socio-economic consequences of this diffusion are well documented. Because mobile adoption functions as a catalyst for market participation, financial inclusion and access to public services, extending connectivity to populations that remain offline is a standing policy priority for governments and international agencies. Planning such extensions, however, requires dependable adoption forecasts. Machine-learning models applied to satellite imagery have therefore been used to predict telecommunications demand and handset penetration in data-scarce regions, providing network operators with an evidence base for capital investment (Oughton and Mathur, 2020). Equally, the value of a connection depends on the ease with which it can be used; studies of keypad geometry for SMS composition demonstrate measurable gains in typing efficiency when key arrangements are optimised for thumb kinematics and for the character frequencies of a given language (Azad et al., 2010; Sharmeen et al., 2010).

The sensing capability of the device has opened a second, less familiar frontier. Because a CMOS image sensor responds to individual photon and charged-particle events, ordinary handsets can be enrolled in distributed scientific campaigns. The Distributed Electronic Cosmic-ray Observatory (DECO), for example, classifies camera frames captured by volunteers worldwide in order to detect cosmic-ray muons and ionising radiation from radioactive decay (Li et al., 2023). In a different modality, dedicated retrieval pipelines convert a photograph of printed sheet music into a query against a symbolic music database, returning the corresponding MIDI passages without any trainable parameters (Tsai et al., 2020; Yang et al., 2020). In both cases the scientific or artistic instrument is a consumer-grade phone already in the user's pocket.

Against this record of utility stands a heterogeneous set of misuses. The first is technical and largely invisible to end users: mobile back ends depend heavily on data-centric programming libraries whose APIs impose subtle structural constraints, and violations of those constraints produce run-time failures, silently incorrect outputs and systemic instability in production analytics (Galappaththi et al., 2024). The second is societal: notwithstanding decades of evidence, a substantial fraction of the public continues to attribute cancer risk to handset radio-frequency emissions. This belief is inconsistent

with elementary physics, since photons in the mobile-communication band carry energies several orders of magnitude below the threshold required to break covalent bonds or ionise DNA (Leikind, 2010; Natarajan, 2012). The two misuses are rarely discussed together, yet both erode the same public trust in the technology.

The problem addressed in this paper is therefore the absence of an integrated account of what mobile devices can do, what can go wrong in the software that services them, and what is wrongly believed about them. Existing treatments are insufficient for two reasons. First, hardware innovations such as sensor repurposing are typically evaluated in isolation from the software back ends that ingest their output, so the dependency of the former on the integrity of the latter remains unexamined. Second, public-health communication about handset radiation seldom invokes the physical principles that settle the question, leaving misconceptions intact and deflecting attention from the genuine engineering risks of the platform. A holistic framework that simultaneously addresses hardware utility, software integrity and scientific literacy is accordingly required.

The contributions of this paper are as follows:

(i) A taxonomy is introduced that organises the novel uses of cell phone hardware into scientific sensing, cross-modal retrieval, adoption forecasting and interaction-design categories, and relates each category to its backend data-processing requirements.

(ii) The misuse landscape is analysed in two registers: data-centric API misuse in mobile analytics software, and the debunked but durable societal narrative connecting handset radiation to carcinogenesis.

(iii) The Mobile Utility and Risk Assessment Framework (MURAF) is proposed as a three-module pipeline integrating deterministic on-device feature extraction, infrastructure and interface mapping, and language-model-assisted misuse detection with educational mitigation.

(iv) A reproducible evaluation protocol with quantified target metrics is specified so that the framework can be validated empirically.

The remainder of the paper is organised as follows. Section 2 reviews the related literature across the three thematic strands and identifies the research gap. Section 3 specifies the MURAF methodology and its evaluation plan. Section 4 discusses implications, limitations, ethical considerations and future directions, and Section 5 concludes.

2. Review of Related Literature

The literature relevant to this study falls into three groups: work that repurposes handset hardware for sensing and cross-modal retrieval; work on network adoption and user-interface optimisation; and work on the misuse of mobile software infrastructure and on societal misconceptions about the technology. Table 1 summarises representative studies from each group.

2.1. Advanced sensing applications and cross-modal retrieval

The first strand examines how native handset components, principally the camera sensor, can be conscripted for analytical tasks. The DECO project distributes an application that collects dark frames and cosmic-ray candidates from volunteer phones; image morphology discriminates particle-induced tracks from sensor artefacts, and convolutional neural networks perform the final classification (Li et al., 2023). The approach converts a globally dispersed population of handsets into a detector array whose geographic coverage no single laboratory could match, and it does so at negligible marginal cost.

A second line of work addresses the optical channel between the physical and the symbolic. Sheet-music retrieval systems accept a photograph of printed notation and return matching passages from a MIDI corpus. Rather than training an end-to-end recogniser, these systems compute a ‘bootleg score’: a compact representation of noteheads relative to staff lines that is invariant to the distortions introduced by handheld photography, and that supports alignment by dynamic programming (Tsai et al., 2020; Yang et al., 2020). The reported F-measure of approximately 0.89 on benchmark photographs indicates that a deterministic pipeline can rival learned alternatives while satisfying the runtime and memory limits of a mobile device.

The strength of this strand is its democratising effect: scientific and musicological tooling becomes available through hardware that users already own. Its weakness is fragility. Deterministic feature extractors are tuned to particular imaging conditions, and performance degrades under poor illumination, motion blur, heavy compression or the wide variation in sensor quality across the installed base. The present study does not propose a new extractor; it situates such extractors within a broader account of device utility and of the backend requirements they imply.

2.2. Telecommunications adoption and interface optimisation

The second strand concerns the reach of the network and the ergonomics of the interface. Oughton and Mathur (2020) demonstrate that convolutional models trained on publicly available satellite imagery can predict cell phone adoption and telecommunications demand at sub-national resolution, offering planners a proxy for costly household surveys. The approach is attractive precisely because it is scalable, yet its reliance on night-time-luminosity and built-environment proxies means that it captures aggregate economic activity more reliably than individual affordability.

Interface research addresses the last centimetre of the network. Azad et al. (2010) and Sharmeen et al. (2010) reorganise the twelve-key layout so that frequently co-occurring characters in Bengali and English are assigned to positions reachable by natural thumb movement, reporting reductions in keystroke time relative to the conventional arrangement. Although software keyboards have since displaced physical keypads on smartphones, the underlying principle—that interaction cost is a function of language statistics and motor anatomy—remains directly applicable to predictive touch interfaces and to the feature-phone segment that still serves a large share of users in low-income regions.

The contribution of this strand is its insistence on accessibility and economic viability as preconditions rather than afterthoughts; its limitation is that adoption forecasts and keypad studies are rarely linked to the applications that connectivity is meant to enable. The framework developed here treats coverage and usability as foundational layers upon which sensing applications can subsequently be deployed.

2.3. System misuses and societal misconceptions

The third strand documents the frictions of the ecosystem. On the technical side, Galappaththi et al. (2024) conduct an empirical study of API misuse in data-centric libraries—the array, dataframe and tensor abstractions on which mobile analytics back ends depend—and show that violations of shape, axis and type constraints are both frequent and consequential, producing run-time exceptions, silent miscomputation and cascading failures. Because such libraries are dynamically typed and highly polymorphic, conventional static analysers recover little of the information needed to detect these violations. A subsequent study therefore evaluates large language models as detectors and repairers, supplying the model with API directives and data-structure summaries extracted from the code base (Galappaththi et al., 2025).

On the societal side, the literature is equally clear. Leikind (2010) and Natarajan (2012) show that the photon energy associated with mobile frequencies is of the order of micro-electron-volts, whereas covalent bond dissociation requires several electron-volts; the photoelectric effect therefore cannot be initiated, and no plausible mechanism exists by which handset radiation could damage DNA. Epidemiological association studies that report elevated risk are, on this account, better explained by recall bias, confounding and low statistical power than by biophysics. Both halves of this strand are rigorous, yet both are communicated in specialist registers that rarely reach either the developer community or the general public. The present paper responds by positioning software API constraints as the genuine technical risk of the mobile era and by contrasting them explicitly with the physically baseless anxieties that dominate public discourse.

Table 1 - Representative literature across the three thematic strands of the mobile telephony ecosystem.

Strand	Representative study	Core contribution	Principal limitation
Sensing and cross-modal retrieval	Li et al. (2023)	Repurposes CMOS camera sensors in volunteer handsets to detect cosmic-ray muons and ionising radiation; convolutional classification of sensor frames.	Sensitivity varies with sensor generation; requires large volunteer base for spatial coverage.
Sensing and cross-modal retrieval	Tsai et al. (2020); Yang et al. (2020)	Deterministic ‘bootleg score’ representation enables retrieval of MIDI passages from photographs of printed sheet music (F-measure ≈ 0.89).	Sensitive to illumination, blur and skew; thresholds require per-device tuning.
Adoption and interface	Oughton and Mathur (2020)	Predicts cell phone adoption and telecom demand from satellite imagery to guide infrastructure investment.	Proxy features such as night-time luminosity misrepresent micro-level affordability.
Adoption and interface	Azad et al. (2010); Sharmeen et al. (2010)	Language-aware keypad layouts reduce SMS entry time by exploiting character frequency and thumb kinematics.	Assumes physical keypads; benefits diminish on touchscreen and predictive-text interfaces.
System misuse and myth	Galappaththi et al. (2024; 2025)	Characterises API misuse in data-centric libraries and evaluates LLM-based detection and repair of constraint violations.	Repair may hallucinate edits; evaluation confined to a limited set of libraries.
System misuse and myth	Leikind (2010); Natarajan (2012)	Shows that mobile-frequency photon energy is far below bond-dissociation thresholds, refuting a radiogenic mechanism for cancer.	Technical framing rarely translated into public-health communication.

2.4. Synthesis and research gap

Read together, the three strands describe a single value chain that the literature habitually examines in segments. Sensing and retrieval studies presume a functioning application back end; back-end studies presume a user population that trusts the technology; adoption studies presume interfaces that the population can actually operate. No existing framework traces this chain end to end, and none weighs the technical and the societal failure modes of mobile telephony on a common footing. The gap this paper addresses is therefore methodological: the absence of an integrated assessment procedure that begins with the physical signal captured by the handset, follows it through the software that processes it, and terminates in the public understanding of the resulting system.

3. Proposed Methodology: The Mobile Utility and Risk Assessment Framework

3.1. Framework overview

The Mobile Utility and Risk Assessment Framework (MURAF) is a three-stage pipeline that evaluates both the functional deployment of mobile applications and the mitigation of the misuses associated with them. Its rationale is to give researchers, developers and regulators a single procedure that maps hardware inputs to software processing while concurrently executing safety and educational checks. The stages are Data Acquisition and Sensing; Infrastructure and Interface Mapping; and Misuse Detection and Educational Mitigation. Figure 1 shows the architecture, including the feedback path by which validated corrections return to the first stage. The design is deliberately modular. Each stage produces a verifiable artefact—a feature vector, a coverage and risk register, or a hardened build with accompanying public guidance—so that the framework can be applied partially, and so that failures can be localised. Table 2 summarises the objectives, techniques and outputs of the three modules.

Table 2 - Composition of the MURAF framework.

Module	Objective	Key techniques	Principal output
I. Data acquisition and sensing	Capture environmental or cross-modal input with unmodified handset hardware and extract compact, discriminative features on the device.	Classical computer vision; staff-line normalisation; image morphology; dynamic-programming alignment; rule-based thresholds.	Mid-level feature vectors and confidence scores, computed without network access.
II. Infrastructure and interface mapping	Establish where connectivity and usable interfaces exist, and inventory the software components that will process the captured data.	Satellite-imagery regression; proxy-feature validation; API-directive and dependency extraction; constraint modelling.	Geospatial coverage forecast and a backend risk register keyed to data-centric libraries.
III. Misuse detection and educational mitigation	Detect and repair API constraint violations, and correct unfounded societal beliefs with evidence-based messaging.	LLM prompting with structural context; automated patch generation; human-in-the-loop review; plain-language physics explainers.	Hardened build, audit trail, and a public scientific-literacy brief.

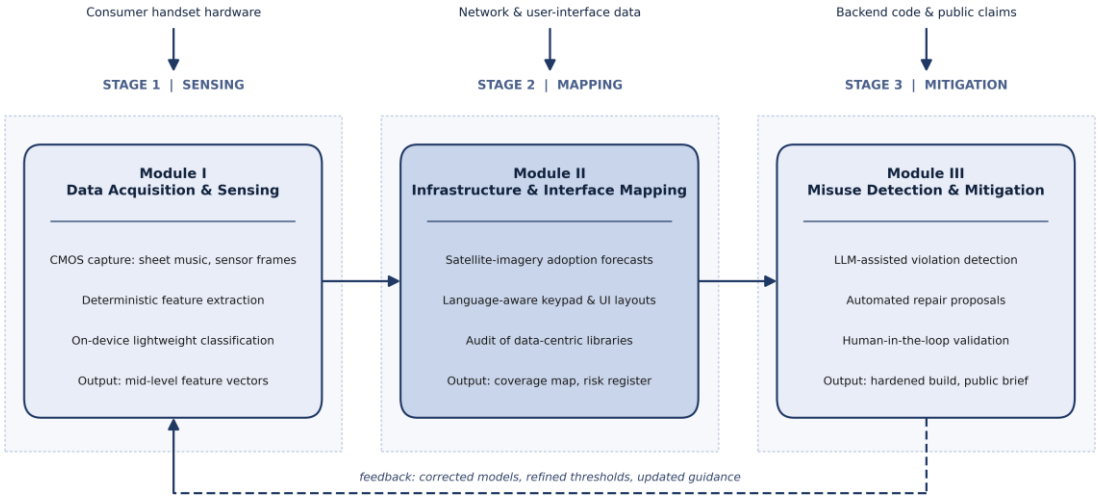


Fig. 1 - Architecture of the Mobile Utility and Risk Assessment Framework (MURAF): three sequential modules with a corrective feedback path.

3.2. Module I: Data acquisition and sensing

The first module formalises the capture of environmental or cross-modal input using standard mobile hardware. Deterministic rules and classical computer-vision operators extract mid-level representations from camera frames: a photograph of sheet music is rectified to the staff coordinate system and reduced to a bootleg score, which is then matched against the symbolic corpus by dynamic programming rather than by a deep network (Tsai et al., 2020; Yang et al., 2020). Analogously, particle events are identified in CMOS frames by localised morphological criteria before any optional cloud-side classification (Li et al., 2023).

The preference for lightweight deterministic extraction is not a rejection of learning but a response to constraint. On-device budgets for energy, memory and latency are severe, and an application that drains the battery or requires continuous uplink will not be adopted. Deterministic pipelines also yield reproducible outputs, which matters when the same evidence must be auditable across heterogeneous devices. A tunable confidence threshold is exposed so that the operating point can be shifted from high recall, as preferred in exploratory citizen-science campaigns, to high precision, as preferred in retrieval applications.

3.3. Module II: Infrastructure and interface mapping

The second module answers two questions that condition everything downstream: where will the application find its users, and through which software will their data be processed? For the first, the module applies satellite-imagery regression to estimate adoption and demand, following Oughton and Mathur (2020), but it additionally requires that the proxy features be validated against any available ground truth before forecasts are released. For the second, the module performs a dependency audit of the backend code base, enumerating the data-centric libraries in use, the API directives they impose, and the data structures flowing between them.

Interface mapping completes the module. Where a deployment targets users who rely on feature phones or on languages with large character sets, layout guidance derived from character frequencies and thumb kinematics is translated into requirements for the predictive-text layer (Azad et al., 2010; Sharmeen et al., 2010). The output of Module II is consequently twofold: a geospatial forecast with stated uncertainty, and a risk register that ranks backend components by their exposure to constraint violations.

3.4. Module III: Misuse detection and educational mitigation

Once data are captured and the software path is known, the third module executes the integrity checks. Because the data-centric nature of mobile analytics is prone to intricate API misuse, the module incorporates an adaptive, language-model-based verification step. API directives and data-structure summaries are extracted from the candidate code and supplied as context to a model that identifies constraint violations before they manifest as production failures (Galappaththi et al., 2024; Galappaththi et al., 2025). An LLM-based detector is preferred to purely syntactic analysis because the violations of interest depend on runtime shapes and axis semantics that static tools cannot resolve.

Automated repair is treated as a proposal, not a decision. Every generated patch is executed against the project's test suite and reviewed by a developer before merge; patches that fail either gate are discarded and logged, which yields a corpus of rejected repairs useful for subsequent prompt refinement. The module closes with the educational component: the same evidence base is re-expressed as plain-language material establishing the energy scales that separate non-ionising radio-frequency exposure from bond-breaking radiation (Leikind, 2010; Natarajan, 2012). Presenting the software audit and the physics explainer as outputs of one module is intentional, since it directs public attention toward the risks that are real.

3.5. Proposed evaluation protocol

The framework is designed to be falsifiable, and Table 3 specifies the protocol by which each component would be assessed. For the utility components, a benchmark of 5,000 handset photographs is envisaged, stratified by device tier, illumination condition and degree of occlusion, together with simulated cosmic-ray artefacts injected into dark frames at known rates. Retrieval quality is measured by the F-measure, with the target set above the 0.89 reported for laboratory-grade pipelines, and classification quality is measured by the area under the precision-recall curve at a fixed false-positive rate.

For the misuse component, 500 documented API misuses are to be injected into an instrumented mobile-analytics backend that depends on widely used data-science libraries. Detection is scored by F1 at the level of the offending statement, and repair is scored by the proportion of proposals that pass the test suite without altering program semantics. A reliability sub-study re-runs the detector on three seeds and reports variance, since non-determinism in generation is itself a deployment risk. Finally, the educational component is evaluated with a pre- and post-intervention instrument measuring the acceptance of the radiation myth, using a control group that receives conventional warning material.

Table 3 - Proposed evaluation protocol and target metrics for MURAF.

Component	Benchmark	Metric	Target
Cross-modal retrieval (Module I)	5,000 handset photographs of printed sheet music, stratified by device tier, illumination and occlusion.	F-measure of correct passage retrieval; median on-device latency.	F-measure ≥ 0.90; latency ≤ 2 s per query.
Ionising-radiation sensing (Module I)	Simulated cosmic-ray and decay events injected into dark frames at known rates across sensor models.	Area under the precision-recall curve; false-positive rate per hour.	AUPRC ≥ 0.85; false positives ≤ 1 per hour.
Adoption forecasting (Module II)	Satellite tiles paired with census- and survey-derived adoption ground truth for two regions.	Mean absolute percentage error; rank correlation with ground truth.	MAPE $\leq 15\%$; Spearman $\rho \geq 0.80$.
API misuse detection and repair (Module III)	500 injected misuses in an instrumented backend using data-centric libraries, with three random seeds.	Statement-level F1; proportion of patches passing the test suite; variance across seeds.	F1 ≥ 0.80; verified repair $\geq 60\%$; F1 variance ≤ 0.03.
Scientific-literacy intervention (Module III)	Pre- and post-intervention survey with treatment and control groups.	Change in acceptance of the radiation–cancer claim; retained knowledge at four weeks.	Relative reduction $\geq 30\%$; retention $\geq 70\%$.

4. Discussion

4.1. Practical implications

For network operators and development agencies, the framework makes the sequence of investment decisions explicit: forecast demand from imagery, validate the forecast against whatever ground truth exists, and only then commit capital. Because adoption forecasts can be produced at sub-national resolution, greenfield areas can be identified with greater confidence than national aggregates permit, which is a precondition for narrowing the digital divide (Oughton and Mathur, 2020). For application developers, the most consequential implication is architectural: lightweight, deterministic retrieval and classification can be executed at the edge, reducing server load, preserving privacy and making capabilities such as instant music transcription available on mid-range hardware (Tsai et al., 2020).

For platform maintainers, the framework reframes API misuse as a first-order reliability concern rather than a coding-style issue. The evidence that constraint violations in data-centric libraries cause silent miscomputation, not merely crashes, implies that mobile analytics outputs cannot be trusted unless the code path that produced them has been audited (Galappaththi et al., 2024). Embedding an LLM-assisted audit in the release pipeline operationalises that requirement at acceptable cost.

4.2. Limitations and failure modes

Four failure modes are anticipated and should be monitored explicitly.

(i) Hardware heterogeneity. Camera sensor quality varies widely across the installed base, and deterministic vision operators calibrated on flagship devices may degrade sharply on low-cost hardware, reducing retrieval accuracy and radiation-detection sensitivity alike.

(ii) Domain shift in adoption forecasting. Models trained on imagery from one region transfer poorly to regions with different infrastructural topography or purchasing behaviour, producing confident but inaccurate demand estimates.

(iii) Repair hallucination. Language models may generate plausible but incorrect patches for data-centric libraries, introducing subtle data corruption in place of the original defect; human review and test gating are therefore mandatory rather than optional.

(iv) Educational backfire. Didactic messaging about radiation can increase salience and, in some audiences, entrench rather than reduce the misconception, which is why the intervention is evaluated against a control group.

A further limitation is structural: the evaluation protocol in Table 3 has not yet been executed at scale, so the present contribution is a methodological proposal supported by secondary evidence rather than a report of completed experiments. The target metrics are therefore stated as thresholds to be tested, not as achieved results.

4.3. Ethical and societal considerations

Two equity concerns arise directly from the analysis. First, the most advanced capabilities described here presume smartphone hardware; populations that rely on basic handsets are not merely underserved by such applications but may be actively excluded from the scientific and cultural participation they enable. Designing deliberately for low-end devices is consequently an ethical requirement, not an optimisation. Second, unfounded anxiety about handset radiation imposes measurable psychological cost and diverts public-health resources from hazards that are real; the failure to communicate elementary physics is itself a public-health failure (Natarajan, 2012). Data-protection obligations form a third consideration: sensing applications that upload camera frames from volunteers must minimise collection, document purpose and provide meaningful consent.

4.4. Future research directions

Three directions follow. The first is the development of fully on-device models that adapt to the optical characteristics of low-end cameras, so that retrieval and detection performance is not a function of purchase price. The second is interdisciplinary work joining software engineering with science communication, in order to design interfaces that explain device mechanics to users while simultaneously protecting them from software vulnerabilities. The third is regulatory: audit evidence of the kind produced by Module III could be incorporated into procurement requirements for public-sector mobile services, converting a voluntary engineering practice into an institutional expectation.

5. Conclusion

The cell phone has completed its transition from a localised communication device into a general platform for computation, scientific sensing and cross-modal retrieval. This paper has surveyed that transition and shown that its most striking instances—cosmic-ray detection with unmodified CMOS sensors, and the retrieval of MIDI passages from photographs of printed music—depend on compact, deterministic representations that respect the energy and latency budgets of consumer hardware. It has also acknowledged the preconditions of such utility: connectivity extended through satellite-imagery-based demand forecasting, and interfaces optimised for the languages and motor anatomies of the populations concerned.

Against this record, the paper has separated two categories of misuse that public debate habitually conflates. The technical category is real and under-reported: violations of structural constraints in the data-centric libraries that service mobile back ends produce silent miscomputation and systemic instability, and they demand the kind of language-model-assisted detection and human-verified repair that the third MURAF module

institutionalises. The societal category is illusory: the attribution of cancer risk to handset radio-frequency emissions is incompatible with photon energetics and the photoelectric effect, and its persistence is a failure of communication rather than of evidence.

The contribution of the paper is to place both categories within a single assessment framework and to attach explicit, falsifiable metrics to each. Realising the potential of the mobile ecosystem will require attention to three things at once: hardware innovation that remains usable on the cheapest devices, software engineering disciplined enough to keep data-centric pipelines correct, and public scientific literacy sufficient to distinguish the risks that exist from those that do not.

REFERENCES

1. Azad, M. A. K., Sharmeen, R., Ahmad, S., & Kamruzzaman, S. M. (2010). Smart Bengali cell phone keypad layout. In Proceedings of the 8th International Conference on Computer and Information Technology (ICCIT 2005), Dhaka, Bangladesh, pp. 1208-1211. arXiv:1009.4979.

2. Galappaththi, A., Nadi, S., & Treude, C. (2024). An empirical study of API misuses of data-centric libraries. Proceedings of the ACM on Software Engineering, 1(FSE). https://doi.org/10.1145/3674805.3686685

3. Galappaththi, A., Ribeiro, F., & Nadi, S. (2025). Detecting and fixing API misuses of data science libraries using large language models. arXiv:2509.25378. https://arxiv.org/pdf/2509.25378

4. Leikind, B. (2010). Do cell phones cause cancer? Skeptic, 15(4). arXiv:1007.4192.

5. Li, R., Pizzuto, A., Vandenbroucke, J., & Mode, B. (2023). Simulation of ionizing radiation in cell phone camera image sensors. arXiv:2311.17253. https://arxiv.org/pdf/2311.17253

6. Natarajan, V. (2012). The myth of cell phone radiation. Resonance, 17(11), 1048-1053. https://arxiv.org/pdf/1211.5203

7. Oughton, E. J., & Mathur, J. (2020). Predicting cell phone adoption metrics using satellite imagery. arXiv:2006.07311. https://arxiv.org/pdf/2006.07311

8. Sharmeen, R., Azad, M. A. K., Ahmad, S., & Kamruzzaman, S. M. (2010). Completely enhanced cell phone keypad. In Proceedings of the International Conference on Information Management and Business (IMB 2005), Shih Chien University, Taiwan, pp. 217-221. arXiv:1009.4980.

9. Tsai, T. J., Yang, D., Shan, M., Tanprasert, T., & Jenrungrot, T. (2020). Using cell phone pictures of sheet music to retrieve MIDI passages. IEEE Transactions on Multimedia, 22(12). https://doi.org/10.1109/TMM.2020.2973831

10. Yang, D., Tanprasert, T., Jenrungrot, T., Shan, M., & Tsai, T. J. (2020). MIDI passage retrieval using cell phone pictures of sheet music. In Proceedings of the 21st International Society for Music Information Retrieval Conference (ISMIR). arXiv:2004.10347.