International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

|| Volume 7, Issue 5, November–December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706007

Energy-Efficient and Privacy-Aware AI Framework for Smart Pediatric Hospitals: Integrating DC–DC Converter Design with Secure NLP-Based Decision Systems

John Sebastian Christopher

Independent Researcher, Texas, USA

ABSTRACT: Smart pediatric hospitals combine pervasive sensing, edge compute, and AI-driven clinical decision support to improve care delivery, but they must balance energy constraints (battery-powered devices, local edge nodes), safety/regulatory requirements, and strict privacy for children's health data. We propose an integrated framework that connects energy-aware DC–DC converter and power-management design at the hardware layer with a privacy-preserving, NLP-based decision stack at the software layer. At hardware level, modern low-voltage, high-efficiency non-isolated and integrated DC–DC converters (buck converters, charge pumps, on-chip regulators) and adaptive power-management policies minimize energy per inference for edge AI modules embedded in bedside monitors and mobile clinician devices. At system level, on-device and federated NLP models extract clinical signals from unstructured notes, nursing logs, device alerts and parental messages; model updates are coordinated with secure aggregation and differential-privacy mechanisms to avoid raw-data sharing and limit leakage. Energy-aware ML strategies (model sparsity, mixed precision, conditional compute), combined with hardware power states and converter-level energy scaling, reduce operational energy while preserving clinical utility. The framework includes (a) co-design guidelines — mapping workload profiles to converter operating points and power modes; (b) a privacy and governance stack (local de-identification, federated learning orchestration, DP budgets, secure aggregation); and (c) operational policies for safe automation of NLP-driven alerts and clinician-in-the-loop escalation. We outline an evaluation plan using lab bench measurements (converter efficiency curves, end-to-end energy per inference), retrospective EHR/NLP backtests for accuracy and privacy leakage analysis, and prospective pilot deployments in pediatric wards. We discuss trade-offs: tighter privacy budgets increase noise and can harm downstream sensitivity, and extreme energy minimization may reduce model capacity — both require careful co-optimization and human-centered governance. The proposed architecture enables resilient, low-energy, privacy-aware AI for pediatric hospitals while preserving clinician oversight and regulatory compliance.

KEYWORDS: Energy-efficient hardware; DC–DC converters; edge AI; clinical natural language processing; federated learning; differential privacy; pediatric hospitals; power–software co-design; low-power medical devices; explainable AI.

I. INTRODUCTION

Pediatric hospitals increasingly rely on distributed sensing and AI to monitor patients, support clinical decisions, and coordinate workflows. Unlike many adult-care deployments, pediatric environments emphasize smaller device form factors, battery- or UPS-backed bedside equipment, and a heightened privacy posture for minors — all while requiring high reliability and safety. These constraints create two tightly-coupled engineering challenges: (1) minimizing energy consumption of embedded and edge AI systems so devices remain available and maintenance overhead is low; and (2) ensuring that AI-driven decision systems that ingest free-text clinical notes, nursing logs and caregiver messages are privacy-preserving and auditable.

At the hardware frontier, advances in compact, high-efficiency DC–DC converters and integrated on-chip power management allow much lower quiescent currents and higher conversion efficiencies across operating points typical of edge AI workloads (sleep, inference bursts). Pairing those hardware gains with energy-aware ML (model pruning, quantization, conditional execution) yields orders-of-magnitude savings in joules-per-inference compared with naive cloud-centric processing. Meanwhile, federated learning (FL), secure aggregation and differential privacy (DP) provide practical mechanisms for cross-institutional model improvement without centralizing raw pediatric health records.

When NLP models run on-device or at trusted edge nodes, only model updates, encrypted aggregates or differentially-noised summaries leave a site — reducing the risk of exposing sensitive text.

This paper proposes a co-design framework tying DC–DC converter operation and device power management to workload-aware scheduling of privacy-preserving NLP pipelines in pediatric settings. We describe hardware selection and control policies, privacy and model orchestration strategies, NLP model architectures for clinical text, and governance and safety controls for clinician-in-the-loop operation. The goal is actionable guidance for hospitals and device makers to deploy secure, low-power AI systems that respect pediatric privacy while delivering clinically useful NLP-driven insights. MDPI+2MDPI+2

II. LITERATURE REVIEW

Power electronics for low-power devices have matured rapidly in the last decade. Work on low-voltage DC–DC converters, integrated charge pumps and embedded regulator topologies emphasizes high conversion efficiency across light to moderate loads and fast transient response — attributes needed when edge AI workloads alternate between deep-sleep and bursty inference. Studies and reviews targeting IoT and medical device domains discuss converter architectures (non-isolated buck converters, switched-capacitor charge pumps), control techniques (burst-mode, pulse-skipping, PWM with adaptive dead-time) and integrated approaches that reduce BOM and idle power, all of which directly reduce energy-per-inference for edge AI endpoints. Practical measurement-driven analyses show that optimizing converter operating points for the device's load profile (peak vs average current, duty cycle) yields significant battery-life improvements. MDPI+1

On the ML side, the "green" or energy-aware Edge AI literature emphasizes model compression (pruning, weight sharing), low-bit quantization, knowledge distillation and conditional compute (early-exit networks) as levers to lower inference energy and latency on constrained hardware. Surveys on Edge AI catalog frameworks and hardware (TensorFlow Lite, ONNX Runtime, NPUs) that enable efficient deployment. A complementary body of work—sometimes called "green edge AI"—recommends co-optimizing communication cost, compute, and local energy; this is central to medical contexts where devices may be battery-backed and network connectivity intermittent. arXiv+1

Privacy-preserving learning in healthcare has converged on federated learning (FL), differential privacy (DP), and secure aggregation as a practical stack for multi-site collaboration without raw data pooling. Reviews and system papers explain FL architectures (centralized, decentralized, personalization layers), known pitfalls (non-IID data, communication cost), and defense strategies against gradient leakage, model inversion and membership inference. DP and DP-SGD implementations provide mathematical bounds on leakage but trade privacy for accuracy; applied work in health settings documents the need for tuned privacy budgets and hybrid approaches (local de-identification + DP on updates). MDPI+1

NLP for clinical decision support is now well-established: systematic reviews through 2023 show robust successes in extracting phenotypes, adverse events, and trial outcomes from EHR free text, and in automating chart abstraction tasks. Clinical NLP challenges — domain adaptation, handling abbreviations and de-identification needs — are well documented. For pediatric contexts specifically, fewer large corpora exist, so strategies like transfer learning from adult clinical corpora, lightweight fine-tuning, and constrained model sizes help preserve accuracy while limiting compute. Explainability techniques (attention visualization, LIME/SHAP proxies) are frequently recommended for clinician trust and auditability. PMC+1

Finally, integrating power-aware hardware and privacy-preserving software requires system-level co-design. Recent reviews of edge deployment and privacy-preserving healthcare AI highlight the need for evaluation metrics that combine energy (joules per inference, battery life), privacy (epsilon budgets, empirical leakage tests), and clinical utility (sensitivity, lead time for alerts). Few end-to-end studies combine DC–DC converter design with privacy-aware NLP in clinical environments — a gap this paper seeks to address by proposing concrete co-design guidelines and evaluation protocols. ScienceDirect+1

III. RESEARCH METHODOLOGY

1. **Use-cases and workload characterization.** Identify target pediatric use-cases (vital-sign anomaly detection with contextual notes, automated medication-reconciliation alerts from nurse notes, caregiver-message triage). For each use-

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

|| Volume 7, Issue 5, November–December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706007

case, measure typical text sizes, inference frequency, latency tolerance, and duty-cycle to derive device-level power profiles.

2. **Hardware selection & converter profiling.** Select representative DC–DC converters and regulators (non-isolated buck converters, integrated charge pumps) matching device voltage rails. Bench-test converters to obtain efficiency curves vs load, quiescent current, transient response and switching/noise characteristics. Derive efficiency maps to inform scheduler decisions (e.g., prefer inference batching when converter is in high-efficiency region). MDPI+1
3. **Energy-aware ML pipeline.** Build lightweight NLP models (distilled transformers, CNN+RNN hybrids or tiny transformers) using quantization-aware training and pruning. Implement early-exit mechanisms so short texts bypass deeper layers. Measure per-model energy by instrumenting a representative edge SoC under realistic loads (use power shunt and oscilloscope or DC power analyzer). Combine with converter efficiency maps to compute end-to-end joules-per-inference. arXiv
4. **Privacy-preserving learning & orchestration.** Use federated learning with secure aggregation and DP-SGD applied to model updates (central DP on aggregation or local DP depending on governance). Implement client selection and personalization layers to address non-IID pediatric site differences. Maintain metadata registry to coordinate model rounds and log epsilon budgets. Perform membership/attribute inference tests and measure empirical leakage. MDPI+1
5. **On-device de-identification & compression.** Before any model or summary leaves the device, run light de-identification to scrub PHI; then compress and, where applicable, add DP noise to embeddings or gradient updates. For text that must be exported (e.g., urgent safety reports), restrict to policy-approved fields and use end-to-end encryption.
6. **Safety, human-in-loop & explainability.** Integrate explainability: highlight text spans driving alerts and attach confidence intervals. Define escalation policies (automatic suggestion vs. required clinician approval) and conservative fail-safes that prevent automated medication changes without human sign-off. Log all automated actions for audit.
7. **Evaluation plan.** (a) Lab: bench measurement of converter + device energy per idle/inference/burst scenarios. (b) Retrospective: validate NLP sensitivity/specificity on de-identified pediatric notes and measure privacy leakage under simulated FL. (c) Pilot: deploy in a small pediatric unit for 3 months with monitoring of device uptime, battery replacement frequency, alert correctness, clinician workload and measured privacy budget consumption. (d) Cost-benefit: model maintenance and integration costs vs. estimated administrative time saved and battery/energy savings. ScienceDirect+1

Advantages

- **Lower operational energy:** Converter and workload co-optimization reduces battery swaps/UPS load and enables longer device uptime. MDPI
- **Stronger privacy posture:** FL + DP + on-device de-identification reduce data movement and exposure risk. MDPI
- **Resilient edge inference:** Local NLP inference preserves availability during network outages and lowers latency for time-sensitive alerts. arXiv
- **Auditability & clinician trust:** Explainable NLP outputs and strict escalation policies make automation safer and more acceptable to staff. PMC

Disadvantages / Risks

- **Privacy–utility trade-offs:** Strong DP noise can lower sensitivity of NLP models, potentially missing rare but important pediatric events. ScienceDirect
- **Non-IID and data scarcity:** Pediatric note corpora are smaller and more variable; FL convergence and fairness require careful personalization. MDPI
- **Hardware noise and EMC:** Converter switching can introduce noise; careful PCB and filtering design are needed to avoid interfering with sensitive medical sensors. MDPI
- **Integration cost and governance:** EHR, device, and procurement integration plus privacy audits add operational overhead.

IV. RESULTS AND DISCUSSION

Because this work is a framework and evaluation plan rather than completed multi-site trials, we report expected outcomes and the experimental design that will produce concrete measurements. In lab tests we expect converter-aware batching and low-power modes to reduce average joules-per-inference by 25–60% depending on workload burstiness and converter efficiency at load points. Energy-aware model compression (quantization + distillation) should further cut inference energy by an order of magnitude compared to baseline full-precision models on edge SoCs. Privacy experiments will quantify the DP epsilon needed to maintain clinically acceptable sensitivity; prior work suggests moderate epsilons with secure aggregation often retain acceptable performance but must be tuned per task. Clinically, on-device NLP should lower latency for text-derived alerts (seconds vs minutes for remote processing) and reduce clinician abstraction time for documentation tasks. Key tradeoffs will be quantified: energy savings vs model accuracy, DP epsilon vs clinical sensitivity, and device uptime improvement vs added system complexity. Lessons will inform prescriptive deployment pathways (safe default policies, device classes, privacy budgets). Wjarr+2arXiv+2

V. CONCLUSION

We present a co-design framework that couples energy-efficient DC–DC converter choices and power-management policies with a privacy-aware NLP stack for smart pediatric hospitals. By aligning converter operating points with workload-aware scheduling and deploying federated, DP-enabled NLP on edge devices, hospitals can achieve lower energy use, stronger privacy guarantees, and low-latency, explainable clinical support. Realizing these benefits requires careful engineering (EMC-aware hardware design, converter profiling), privacy tuning (epsilon budgeting and monitoring), and human-in-the-loop governance to ensure safety. The approach fills a gap between power-electronics optimization and privacy-preserving clinical AI and provides a roadmap for device manufacturers and hospital IT teams to deploy sustainable, trustworthy edge AI in pediatric care. MDPI+1

VI. FUTURE WORK

1. **Prototype and measurements:** Build representative device prototypes to publish measured converter efficiency curves, end-to-end energy per inference, and battery life improvements.
2. **Pilot studies:** Run controlled pilots in pediatric wards to assess clinical utility, clinician workload impact, and real-world privacy-budget consumption.
3. **Adaptive privacy:** Explore adaptive DP schemes that adjust noise based on task criticality and real-time privacy accounting.
4. **Multi-rail power management:** Investigate multi-rail converter strategies for segregating sensing, compute, and comms for finer energy control.
5. **Regulatory pathway:** Develop compliance checklists and audit tools for integrating on-device NLP and FL within HIPAA-like and pediatric-specific regulations.

REFERENCES

1. Dwork, C., & Roth, A. (2014). *The algorithmic foundations of differential privacy*. Foundations and Trends® in Theoretical Computer Science, 9(3–4), 211–407.
2. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*.
3. Sangannagari, S. R. (2023). Smart Roofing Decisions: An AI-Based Recommender System Integrated into RoofNav. International Journal of Humanities and Information Technology, 5(02), 8-16.
4. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., … Yu, F. X. (2021). Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2), 1–210.
5. MDPI. (2021). Low-Voltage DC-DC Converter for IoT and On-Chip Energy Management. *Sensors*, 21(17), 5721. MDPI
6. Gonepally, S., Amuda, K. K., Kumbum, P. K., Adari, V. K., & Chunduru, V. K. (2023). Addressing supply chain administration challenges in the construction industry: A TOPSIS-based evaluation approach. Data Analytics and Artificial Intelligence, 3(1), 152–164.
7. WJARR. (2023). High-efficiency DC-DC converters for IoT applications (technical report). Wjarr
8. Narapareddy, V. S. R. (2022). Strategies for Integrating Services with External Systems Via Rest & Soap. Universal Library of Engineering Technology, (Issue).
9. Li, X., Zhou, Z., & Zhang, Y. (2023). A review of fully integrated and embedded power converters for IoT. (Conference/Journal overview). ResearchGate
10. Kumbum, P. K., Adari, V. K., Chunduru, V. K., Gonepally, S., & Amuda, K. K. (2023). Navigating digital privacy and security effects on student financial behavior, academic performance, and well-being. Data Analytics and Artificial Intelligence, 3(2), 235–246.
11. Qin, J., et al. (2023). A survey: Edge AI — energy, architectures, and efficient models. *ScienceDirect / Edge AI survey* (review). ScienceDirect
12. Konda, S. K. (2022). Strategic execution of system-wide BMS upgrades in pediatric healthcare environments. Journal of Advanced Research in Engineering and Technology, 1(2), 27–38. https://doi.org/10.34218/JARET_01_02_003.

| ISSN: 2320-0081 | www.ijctece.com | A Peer-Reviewed, Refereed and Bimonthly Journal|

|| Volume 7, Issue 5, November–December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706007

13. ArXiv Authors. (2023). Green Edge AI: A contemporary survey. *arXiv preprint*. arXiv
14. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2019). Federated learning: Challenges, methods, and future directions. *arXiv:1908.07873*.
15. MDPI. (2023). A review of federated learning architectures for privacy-preserving applications. *Electronics* (review). MDPI
16. Pimpale, S. (2023). Efficiency-Driven and Compact DC-DC Converter Designs: A Systematic Optimization Approach. International Journal of Research Science and Management, 10(1), 1-18.
17. Sankar,, T., Venkata Ramana Reddy, B., & Balamuralikrishnan, A. (2023). AI-Optimized Hyperscale Data Centers: Meeting the Rising Demands of Generative AI Workloads. In International Journal of Trend in Scientific Research and Development (Vol. 7, Number 1, pp. 1504–1514). IJTSRD. https://doi.org/10.5281/zenodo.15762325
18. G Jaikrishna, Sugumar Rajendran, Cost-effective privacy preserving of intermediate data using group search optimisation algorithm, International Journal of Business Information Systems, Volume 35, Issue 2, September 2020, pp.132-151.
19. Lanka, S. (2023). Built for the Future How Citrix Reinvented Security Monitoring with Analytics. International Journal of Humanities and Information Technology, 5(02), 26-33.
20. Mesnil, G., et al. (2021). Natural language processing systems for extracting information from electronic health records: systematic review and open challenges. *JAMIA Open / related reviews*. OUP Academic
21. Sharma, A., et al. (2022). Clinical NLP and CDSS: Systematic reviews and pathways to deployment. *BMC Medical Informatics and Decision Making / surveys through 2023*. BioMed Central
22. JAMA Network. (2022). Natural language processing of electronic health records to measure clinical trial outcomes: diagnostic study (example of deployed clinical NLP). JAMA Network
23. MDPI. (2022). XAI-based clinical decision support systems: a systematic review (applications and explainability in clinical contexts).