

# BRIE-RAM: Exploring Biodegradable Processing-Using-Memory for Sustainable Edge Sensing

Aniket Das

Ryan Wong

Arjun Tyagi

Saugata Ghose

University of Illinois Urbana-Champaign

## Abstract

*Smart sensors at the edge have the potential to become ubiquitous, but for them to be effective, we must both (1) meet stringent energy constraints while processing increasingly large amounts of data, and (2) design devices that can sit for years in an environment or on organisms without contributing to e-waste. To address both concerns, we explore the feasibility of BRIE-RAM, which performs processing-using-memory inside biodegradable resistive RAM. We evaluate BRIE-RAM’s capabilities and shortcomings through a contemporary lens, and then discuss future research needs and ethical issues to tackle in order to bring BRIE-RAM to fruition.*

## 1. Introduction

Modern edge computing paradigms are ushering in new capabilities and use cases for *smart sensing*, where we potentially embed edge devices integrated with environmental sensors to monitor a wide range of conditions and phenomena. Applications for smart sensing are very diverse, ranging from wildlife tracking [8, 35] to crop and forest management [14, 49] to traffic planning [32] to disaster warnings [1]. With the advent of modern radio technologies such as long-range wide-area networks (LoRaWAN) [25], smart sensors can collect data, coordinate with each other, and report back data with modest networking costs.

However, a key obstacle to smart sensing remains the tight energy constraints at the edge. A commonplace vision for smart sensing is to install edge sensor devices ubiquitously, and then to leave the devices at their location (or on their animal) for years at a time, requiring ultra high energy efficiency to survive off of battery power. Even with technologies such as LoRaWAN, a pure cloud-based computational model requires excessive data to be transmitted to the cloud, incurring significant battery drain due to networking energy or severely constraining the amount and type of data that can be collected [9]. Unfortunately, while local data processing (e.g., filtering, event detection) can eliminate the networking bottleneck, conventional von Neumann architectures (which have to date been the design of choice for smart sensors) incur significant inefficiencies whenever they perform such processing due to the high energy and performance costs of CPU-memory data movement, which are orders of magnitude larger than CPU computation [2, 6].

An emerging solution for practical smart sensing is the use of *processing-using-memory* (PUM) [41]. PUM is a non-von-Neumann paradigm that exploits the electrical characteristics of memory devices to perform computation directly within the devices, allowing it to eliminate data movement overheads across a wide range of data-intensive applications [15, 37, 45, 46, 51]. Effectively, across large memory chips, PUM exposes potentially millions of ways of parallelism, which can be well

sued to many of the parallel data analysis techniques used for smart sensing.

While PUM offers a potential solution to the strict energy constraints of smart sensing, existing PUM works and prototypes are made from inorganic materials such as silicon, using forever chemicals, resulting in harmful e-waste both during production and for chips left in their environment after their lifetimes expire [11, 22]. Such e-waste is particularly troubling for smart edge sensors, as we envision them being deployed in various settings from urban sewage systems to remote rainforests, with few and often difficult avenues for reclamation. Specifically, these devices are more easily lost (e.g., forgotten, washed away, moved or consumed by wildlife) compared to their warehoused counterparts. With the number of wide-area devices expected to cross 5 billion in 2026 and 8 billion by 2031 [44], expanding the reach of ubiquitous computing threatens to expand the reach of ubiquitous e-waste. Collectively, the way we currently design, maintain, and discard smart sensors promises to present a critical sustainability problem for the future of computing.

Recent research on biodegradable memory devices offers a potential path to sustainable PUM for smart sensing. As ReRAM (resistive metal-oxide RAM) is a popular substrate for silicon-based PUM [28, 31, 46, 47, 51], we examine whether *biodegradable ReRAM* (bio-ReRAM) [38] devices are capable of performing PUM. In this work, we explore three critical questions about the potential of *BRIE-RAM* (*Biodegradable Reasoning Inside Edge RAM*). First, we perform a preliminary analysis of the state of multiple types of bio-ReRAM, and whether their characteristics (e.g., I-V response, programmability, ability to hold state) are amenable to PUM logic. Second, we examine what future bio-ReRAM materials, device, and architecture research should focus on to improve PUM suitability. Third, we discuss potentially emergent ethical issues that arise from using bio-ReRAM.

Our initial studies reveal that egg-based bio-ReRAM [4, 54] exhibits characteristics most amenable for use in BRIE-RAM. Our work reveals a challenging yet promising path towards realizing the potential of BRIE-RAM as the foundations of a sustainable smart sensor architecture.

## 2. Background: Processing-Using-Memory

To reduce data movement overheads, researchers have explored processing-using-memory (PUM), which can provide both orders of magnitude performance and energy benefits [15, 40, 42, 46] compared to conventional von Neumann architectures. The key mechanism behind PUM is its ability to utilize the inherent electrical characteristics of the memory devices to perform computation without the need to move the data to a separate processing element.

Resistive RAM (ReRAM) is a common PUM-enabling memory technology that leverages resistance state to hold data.

The general structure of ReRAM consists of a functional material (e.g., TaOx) sandwiched between two electrodes. To operate the ReRAM device, a set/reset voltage ( $V_{set}/V_{reset}$ ) is applied to the device, which will cause it to switch from a low-resistance state (LRS) to a high-resistance state (HRS) or vice versa, respectively. One recent PUM architecture, RACER [46, 47], enables NOR- and OR-based Boolean operations using ReRAM devices. By applying a specific voltage (e.g.,  $V_{NOR}$ ) to two columns in a memory array that each store input vectors of data, the current along each row applies voltages to the devices in a third, grounded column. This results in the third column storing an output vector that represents the bitwise NOR of the two input vectors. Furthermore, by leveraging the array structure of the memory to process large vectors in parallel, RACER is able to achieve substantial performance and energy benefits. These benefits suggest that PUM architectures may be strong contenders for environmentally friendly smart-sensor applications, which are constrained by severe energy limits.

### 3. BRIE-RAM Requirements

We use the constraints of edge computing and PUM to establish basic goals for BRIE-RAM to be deployed at scale.

#### 3.1. Application Goals

Our first architectural goal is minimal energy consumption. This is one of the most important targets for edge sensor devices. Sensors deployed in remote areas must either rely on a single battery charge or energy harvesting sources such as light [5, 21] or radio waves (RF) [39]. As transmitting data is one of the most energy-intensive tasks for a sensor, smart edge sensors can significantly improve energy consumption by using low-energy local computation to significantly reduce the amount of data that needs to be transmitted across the network. For example, image sensor devices can use local computation to identify and filter out uninteresting images, preventing them from being transmitted.

Our second architectural goal is tolerable application latency. Several smart sensor applications require low latency, such as wildlife tracking sensors designed to provide real-time alerts when humans or animals approach each other’s territories, as conflict between them may be imminent [8]. Others are designed to monitor the biodiversity of a region over time, classifying different observed species [35]. These may require only periodic updates and therefore can tolerate relaxed latency constraints. For edge applications such as filtering useful images, the latency of computation can determine the rate of image capture, as the computation occupies too many resources (e.g., power, energy, memory) to capture and buffer new images. BRIE-RAM should at least be able to support periodic edge computing tasks without straining these resources, though ideally with a roadmap to supporting more latency-critical use cases.

#### 3.2. Reliability Goals

Beyond its energy constraints, BRIE-RAM must be able to have its bio-ReRAM devices operate correctly for the entire desired lifetime. Our first reliability goal is the *bit error rate* (BER) of BRIE-RAM. This is the rate of incorrect output bits when performing PUM operations. The main source of BER we examine is device-to-device variation in the voltages that

must be applied to a ReRAM cell in order for it to switch from the HRS to the LRS or vice versa. Since ReRAM array peripherals apply a single voltage to the entire input and output column of an array when performing PUM operations, the resulting voltage applied to the output device may cause it to incorrectly switch or fail to switch. While PUM is an emerging technology that does not yet have codified standards for BER, we can use the JEDEC standard of  $10^{-15}$  uncorrectable errors in data storage [23] as a target for BRIE-RAM to achieve reliable computation.

BRIE-RAM-based smart sensors should be capable of operating with no maintenance for years or potentially decades, making our second reliability goal the longevity of bio-ReRAM devices. We use two metrics to evaluate this: (1) *retention*, the amount of time that data can be held in a memory cell, which is effectively a measure of the device’s non-volatility; (2) *write endurance*, the number of cycles for which a memory cell can be reset and then set before data can no longer be reliably read or written. Conventional ReRAM has shown estimated retention of up to 10 years [43] and write endurance of  $10^{12}$  cycles [33].

### 4. BRIE-RAM Evaluation & Future Targets

We investigate the viability of deploying BRIE-RAM for real applications based on prior work using three distinct bio-ReRAM device substrates: (1) *Egg*, which uses a thermal-baked egg albumen (i.e., egg white) functional material [4]; (2) *Fiber*, which uses a cellulose nanofiber paper functional material [36]; and (3) *Crystal*, which uses a cellulose nanocrystal functional material [20]. These substrates were chosen based on available data regarding device characteristics and variation, though we note that bio-ReRAM devices can be made with a range of other materials [24, 26, 38, 50].

#### 4.1. Methodology

To evaluate how a BRIE-RAM architecture would operate, we develop a detailed bio-ReRAM array-level model that captures non-idealities in devices and circuitry. Our array level model takes six key parameters that directly affect the behavior of the PUM primitive: (1) the array size; (2) the on/off ratio, or the ratio between the device’s HRS and its LRS,<sup>1</sup> (3) the cycle time, which determines the frequency; (4) the voltages required to set ( $V_{set}$ ) and reset ( $V_{reset}$ ) the devices; (5) the variation in voltage required to induce switching (i.e., the mean  $\mu$  and standard deviation  $\sigma$  of observed necessary threshold voltages); and (6) the *coefficient of variation* (CV), the ratio of the standard deviation to the mean, where a higher CV indicates a broader distribution. Table 1 shows these parameters for the evaluated devices.

Note that as the array size increases, the magnitude and duration of voltage pulses must be adjusted to correctly account for IR (current/resistance; i.e., Ohmic) drops along the array’s wires. In our model, we account for these behaviors, and empirically determine the proper input voltages, while still modelling the array-level behaviors (e.g., IR drop). We use these parameters to simulate BRIE-RAM arrays of size  $32 \times 32$  and perform PUM operations using the RACER architecture [46].

<sup>1</sup> A higher on/off ratio indicates greater resistance to noise.

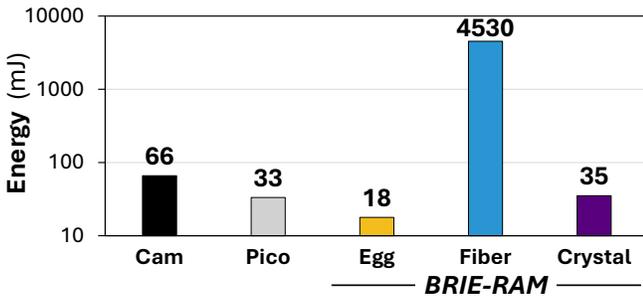
**Table 1: Simulated BRIE-RAM parameters.**

Device	On/Off Ratio	$V_{set}$ / $V_{reset}$	Freq. (kHz)	$\mu$ (set/res.)	$\sigma$ (set/res.)	CV (set/res.)
Egg [4]	$\sim 10^3$	-1 / 5	$10^4$ [54]	-0.55 / 2.6	0.10 / 0.33	-0.17 / 0.13
Fiber [36]	$\sim 10^5$	1 / -1	1 [52]	0.28 / -0.22	0.26 / 0.56	0.93 / -2.55
Crystal [20]	$\sim 10^4$	-1 / 1	1 [52]	-0.19 / 0.21	0.02 / 0.09	-0.10 / 0.41

#### 4.2. Application Challenges

As edge computing applications are moving towards on-device AI [34], we investigate whether BRIE-RAM can support such workloads. We estimate the cost of image classification using the LeNet-5 deep neural network (DNN) [29]. Since LeNet has a smaller memory footprint and simpler computation than more recent DNN architectures such as AlexNet [27] and ResNet [18], it is a good fit for lightweight edge computing tasks [7, 16]. We use reported metrics from two baseline edge computing devices for our inference evaluation: (1) Camaroptera (*Cam*) [9], an energy-harvesting image sensing device that uses a TI MSP430 microcontroller to filter captured frames with a LeNet-structured DNN; and (2) a Raspberry Pi Pico (*Pico*), using an Arm Cortex-M0+ CPU [10].

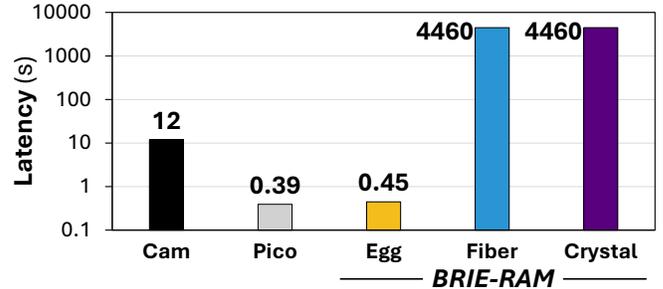
Figure 1 shows the energy consumption of a single DNN inference on each system. *Egg* BRIE-RAM has the lowest energy usage at 0.27 $\times$  and 0.54 $\times$  that of the *Cam* and *Pico* baselines, respectively. *Crystal* BRIE-RAM’s energy usage is 0.53 $\times$  that of *Cam* and on par with *Pico*. This implies that image filtering at the edge with BRIE-RAM can process multiple captured images at once while remaining within the energy budget of prior work. We are currently investigating the sources of variation in energy consumption between bio-ReRAM devices. Further optimization and co-design of DNNs and BRIE-RAM architectures will be necessary to fully leverage the parallelism of PUM while remaining within acceptable energy consumption limits.



**Figure 1: Energy used for a single DNN inference.**

Figure 2 shows the latency of a single DNN inference. The extremely slow switching speed of cellulose devices affects both *Fiber* and *Crystal* equally in terms of performance, resulting in a projected inference time of 74 minutes. While periodic tasks may be able to tolerate longer latencies, such latencies are intolerable for image filtering, and would require massive amounts of memory for buffering captured images and providing the necessary parallelism to clear the buffers. *Egg*, on the other hand, significantly improves upon *Cam*’s

inference latency, and approaches that of *Pico* even with a 10 MHz frequency. As the application pipeline used by *Cam* cannot capture a new image until the current one has been processed, its 12-second inference time causes it to miss many images of interest [9]. Deploying *Egg* BRIE-RAM may even prevent inference from being a bottleneck, as its latency beats Camaroptera’s reported LoRaWAN transmission latency of 0.8 s. We note that BRIE-RAM has significantly greater opportunities to exploit data parallelism compared to *Pico*, opening the path to further gains in application performance. Future research will need to balance these gains with the pressure that increased image processing may place on the system’s energy budget.



**Figure 2: Latency of a single DNN inference.**

While such improvements are not likely to be straightforward and may run into fundamental material limitations, this analysis shows that BRIE-RAM does not have to reach the nanoscale metrics of its inorganic counterparts to be useful.

#### 4.3. Reliability Challenges

We survey prior work to examine if BRIE-RAM can meet the reliability constraints necessary for PUM at the edge.

**Bit Error Rate.** Although prior work has demonstrated that bio-ReRAM can be used to realize Boolean primitives [50], the work neglects to consider array-level implications (e.g., non-idealities in crossbar wiring and peripheral circuitry), limiting their work to small-scale proofs of concept. Using our bio-ReRAM array model, we examine the BRIE-RAM bit error rate (BER) during execution of repeated bitwise OSCAR [47] NOR operations over a period of  $10^5$  cycles. While some bio-ReRAM substrates have displayed on/off ratios approaching that of conventional ReRAM, from Figure 3 we observe significant impacts from device-to-device variation. *Crystal* has the lowest BER at 12%, while *Fiber* has a nearly 50% error rate. This can be explained by *Fiber*’s high coefficient of variation (CV) for  $V_{reset}$ , which means that 34.7% of the time *Fiber* devices require a positive  $V_{reset}$  instead of the negative voltage applied to switch the resistance state. Unfortunately, increasing the magnitude of the applied voltages would not be able to solve this problem due to the mismatch in polarity. Furthermore, dynamically applying a different polarity of  $V_{reset}$  is incompatible with many PUM architectures, which often apply the same voltages across entire wordlines or bitlines to achieve parallel computation.

While lower set and reset voltages can help reduce the energy necessary to carry out operations in BRIE-RAM, these must come with much tighter ranges of variability. *Crystal* uses the same magnitude of voltages, but with its low

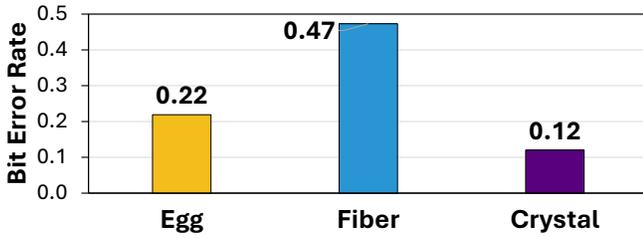


Figure 3: Bit error rate for NOR with bio-ReRAM.

CV, only 0.7% of  $V_{reset}$  thresholds are of the opposite polarity of  $V_{reset}$ . Using these rates of polarity mismatches as a design heuristic can motivate device research to bring bit errors down to the point that the architecture and software can use approximate techniques for bio-ReRAM PUM systems, based on those for conventional ReRAM [17, 19]. This device-to-device variation is perhaps the largest roadblock in the way of successfully deploying BRIE-RAM, and must be addressed by collaboration across the computing stack.

**Longevity.** Figures 4 and 5 show results characterizing write endurance and data retention, respectively, using data reported by prior works [4, 20, 36]. Once again, we observe that the BRIE-RAM results are far below those currently achievable by conventional ReRAM. However, this is a consequence of the methodology used by prior work: they show that the devices remain in good condition with stable resistance states following the reported cycling or retention time, but do not report the point at which writes and reads fail. Therefore, we are unable to draw meaningful conclusions regarding the longevity of a BRIE-RAM architecture. Future characterization of bio-ReRAM devices must test them to failure so an upper bound of reliable operation can be established. This will enable architects to understand the true viability of BRIE-RAM for long-term edge sensing applications.

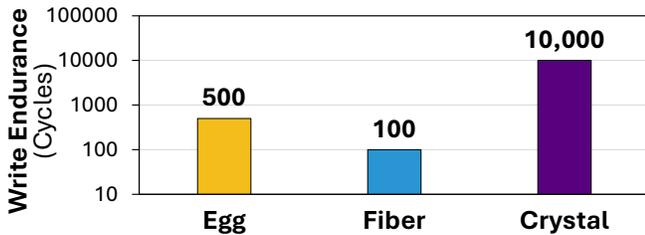


Figure 4: Write endurance of bio-ReRAM devices.

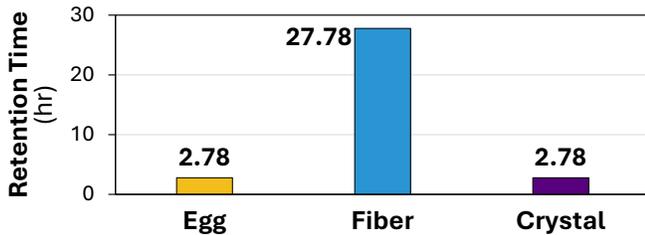


Figure 5: Data retention of bio-ReRAM devices.

## 5. Ethical Challenges of BRIE-RAM

Although BRIE-RAM has the potential to enable more environmentally friendly edge computing, we cannot take for granted that it will be ethical or sustainable by sole virtue of its biodegradability. In addition to the technical challenges of

bio-ReRAM, we consider three ethical challenges in making BRIE-RAM a reality.

First, prior work has made bio-ReRAM devices using eggs [4], gelatin [3], and even salmon DNA [24], and our work suggests that eggs may be a promising candidate for a real BRIE-RAM architecture. However, the use of animal products as the primary component in electronics may conflict with the beliefs of different people and communities. Based on this, we see an opportunity to use this dilemma to motivate further research into plant-based bio-ReRAM.

Second, if BRIE-RAM is able to subsume silicon-based edge sensors, billions of devices will need to be manufactured to meet projected demand [44]. This requires scalable means of sourcing the necessary materials, and manufacturers may find it tempting to turn to factory farms as a solution. In addition to animal treatment concerns, factory farming induces greater natural resource usage and carbon emissions whether they are supplying animal- or plant-based materials. Depending on how lucrative BRIE-RAM may become, manufacturing deals with farms for materials such as eggs may result in consumer shortages and price hikes.<sup>2</sup> As an alternative, BRIE-RAM could spur the establishment of recycling programs, such as using eggs past their sell-by date at supermarkets, or excess/byproduct materials from paper manufacturers.

Beyond the substrate materials themselves, we need to consider ethical issues related to fabrication. For example, one proposed gelatin bio-ReRAM device makes use of embedded cobalt to improve its endurance [30], while a bio-ReRAM device made from silkworm silk uses gold bottom electrodes [53]. Many bio-ReRAM devices use cobalt-doped indium tin oxide electrodes [3, 4, 30]. Ethical concerns have been raised regarding cobalt mining practices (e.g., labor conditions, pollution of the environment and indigenous communities [13]). Gold and tin are also designated as conflict materials that come from regions with high levels of instability and human rights violations [12, 48]. While bio-ReRAM research is still exploring the best materials to use in devices and architectures, With bio-ReRAM materials research still emerging, we should strive to (1) minimize the use of conflict minerals and materials and (2) design devices and architectures with end-to-end biodegradability, requiring effort on many fronts beyond just the computer architecture.

## 6. Conclusion

From our study, we find that BRIE-RAM shows promise to improve the efficiency of local compute in smart edge sensors, while mitigating e-waste. For BRIE-RAM to become a reality, we must realize a number of calls to action. First, there is a need for detailed bio-ReRAM device characterization, with a focus on identifying true device longevity. Second, we should start planning for sustainable materials sourcing, which may require significant investment in improving the traits of non-animal-product bio-ReRAM devices. Third, we need to address significant reliability challenges, which we expect will be exacerbated as we strive to step away from cobalt and conflict materials. Overall, we envision a path to fruition for BRIE-RAM that can spur a new body of research and address emergent ethical and sustainability issues.

<sup>2</sup> In recent years, both the egg and DRAM markets have experienced such shortage-induced price hikes.

## References

- [1] M. H. bin Azami, N. C. Orger, V. H. Schulz, T. Oshiro, and M. Cho, "Earth Observation Mission of a 6U CubeSat With a 5-Meter Resolution for Wildlife Image Classification Using Convolutional Neural Network Approach," *Remote Sens.*, Apr. 2022.
- [2] A. Boroumand *et al.*, "Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks," in *ASPLOS*, 2018.
- [3] Y.-C. Chang and Y.-H. Wang, "Solution-Processed Al-Chelated Gelatin for Highly Transparent Non-Volatile Memory Applications," *Appl. Phys. Lett.*, Mar. 2015.
- [4] Y.-C. Chen, H.-C. Yu, C.-Y. Huang, W.-L. Chung, S.-L. Wu, and Y.-K. Su, "Nonvolatile Bio-Memristor Fabricated With Egg Albumen Film," *Scientific Reports*, May 2015.
- [5] A. Colin, E. Ruppel, and B. Lucia, "A Reconfigurable Energy Storage Architecture for Energy-Harvesting Devices," in *ASPLOS*, 2018.
- [6] B. Dally, "Challenges for Future Computing Systems," 2015, Keynote talk at HiPEAC.
- [7] M. de Prado, M. Rusci, R. Donze, A. Capotondi, S. Monnerat, L. Benini, and N. Pazos, "Robustifying the Deployment of TinyML Models for Autonomous Mini-Vehicles," *Sensors*, Feb. 2021.
- [8] J. S. Dertien *et al.*, "Mitigating Human-Wildlife Conflict and Monitoring Endangered Tigers Using a Real-Time Camera-Based Alert System," *BioScience*, Oct. 2023.
- [9] H. Desai, M. Nardello, D. Brunelli, and B. Lucia, "Camaroptera: A Long-Range Image Sensor With Local Inference for Remote Sensing Applications," *TECS*, May 2022.
- [10] M. Deutel, P. Woller, C. Mutschler, and J. Teich, "Energy-Efficient Deployment of Deep Learning Applications on Cortex-M Based Microcontrollers Using Deep Compression," in *MBMV*, 2023.
- [11] M. Elgamal, A. Mahmoud, G.-Y. Wei, D. Brooks, and G. Hills, "The Environmental Impact of Forever Chemicals in Computing Systems," in *HotEthics*, 2024.
- [12] European Commission, "Conflict Minerals Regulation: The Regulation Explained," [https://policy.trade.ec.europa.eu/development-and-sustainability/conflict-minerals-regulation/regulation-explained\\_en](https://policy.trade.ec.europa.eu/development-and-sustainability/conflict-minerals-regulation/regulation-explained_en), 2021.
- [13] T. C. Frankel, "The Cobalt Pipeline: Tracing the Path From Deadly Hand-Dug Mines in Congo to Consumers' Phones and Laptops," *The Washington Post*, Sep. 2016.
- [14] F. Fuentes-Peñaillillo, K. Gutter, R. Vega, and G. C. Silva, "Transformative Technologies in Digital Agriculture: Leveraging Internet of Things, Remote Sensing, and Artificial Intelligence for Smart Crop Management," *J. Sens. Actuator Netw.*, Jul. 2024.
- [15] D. Fujiki, S. Mahlke, and R. Das, "Duality Cache for Data Parallel Acceleration," in *ISCA*, 2019.
- [16] G. Gobieski, B. Lucia, and N. Beckmann, "Intelligence Beyond the Edge: Inference on Intermittent Embedded Systems," in *ASPLOS*, 2019.
- [17] A. Grossi *et al.*, "Resistive RAM Endurance: Array-Level Characterization and Correction Techniques Targeting Deep Learning Applications," *TED*, Feb. 2019.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *CVPR*, 2016.
- [19] T. Hirtzlin, M. Bocquet, J.-O. Klein, E. Nowak, E. Vianello, J.-M. Portal, and D. Querlioz, "Outstanding Bit Error Tolerance of Resistive RAM-Based Binarized Neural Networks," in *AICAS*, 2019.
- [20] T. Hussain *et al.*, "Cellulose Nanocrystal Based Bio-Memristor as a Green Artificial Synaptic Device for Neuromorphic Computing Applications," *Adv. Mater. Technol.*, Feb. 2022.
- [21] N. Jackson, J. Adkins, and P. Dutta, "Capacity Over Capacitance for Reliable Energy Harvesting Sensors," in *IPSN*, 2019.
- [22] P. Jain, C. Pemberton, I. Hernandez, M. Rosillo, S. West, J. Balkind, and T. Sherwood, "Electronic Waste Footprint of Computer Systems," in *HotEthics*, 2024.
- [23] JEDEC Solid State Technology Assn., "JEP122H: Failure Mechanisms and Models for Semiconductor Devices," Sept. 2016.
- [24] H.-Y. Jeng, T.-C. Yang, L. Yang, J. G. Grote, H.-L. Chen, and Y.-C. Hung, "Non-Volatile Resistive Memory Devices Based on Solution-Processed Natural DNA Biomaterial," *Organic Electronics*, Mar. 2018.
- [25] M. Jouhari, N. Saeed, M.-S. Alouini, and E. M. Amhoud, "A Survey on Scalable LoRaWAN for Massive IoT: Recent Advances, Potentials, and Challenges," *IEEE Commun. Surveys & Tutorials*, Jul.-Sept. 2023.
- [26] G. Kook, S. Jeong, M. K. Kim, S. Lee, N. Choi, and H. J. Lee, "Fabrication of Highly Dense Silk Fibroin Biomemristor Array and Its Resistive Switching Characteristics," *Adv. Mater. Technol.*, Apr. 2020.
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification With Deep Convolutional Neural Networks," in *NIPS*, 2012.
- [28] S. Kvatinsky, D. Belousov, S. Liman, G. Satat, N. Wald, E. G. Friedman, A. Kolodny, and U. C. Weiser, "MAGIC: Memristor-Aided Logic," *TCAS II*, Sep. 2014.
- [29] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-Based Learning Applied to Document Recognition," *Proc. IEEE*, Nov. 1998.
- [30] C.-J. Lee, Y.-C. Chang, L.-W. Wang, and Y.-H. Wang, "Nonvolatile Resistive Switching Memory Utilizing Cobalt Embedded in Gelatin," *Materials*, Jan. 2017.
- [31] H. Li, H. Jin, L. Zheng, and X. Liao, "ReSQM: Accelerating Database Operations Using ReRAM-Based Content Addressable Memory," *TCAD*, Nov. 2020.
- [32] Y. Liu, J. Q. Yu, J. Kang, D. Niyato, and S. Zhang, "Privacy-Preserving Traffic Flow Prediction: A Federated Learning Approach," *IEEE Internet Things J.*, Aug. 2020.
- [33] Q. Luo *et al.*, "Nb<sub>1-x</sub>O<sub>2</sub> Based Universal Selector With Ultra-High Endurance (>10<sup>12</sup>), High Speed (10 ns) and Excellent V<sub>th</sub> Stability," in *VLSIT*, 2019.
- [34] Y. Mao, X. Yu, K. Huang, Y.-J. A. Zhang, and J. Zhang, "Green Edge AI: A Contemporary Survey," *Proc. IEEE*, Jul. 2024.
- [35] J. Müller *et al.*, "Soundscapes and Deep Learning Enable Tracking Biodiversity Recovery in Tropical Forests," *Nature Commun.*, Oct. 2023.
- [36] K. Nagashima *et al.*, "Cellulose Nanofiber Paper as an Ultra Flexible Nonvolatile Memory," *Scientific Reports*, Jul. 2014.
- [37] G. F. Oliveira, A. Olgun, A. G. Yağlıkçı, F. N. Bostancı, J. Gómez-Luna, S. Ghose, and O. Mutlu, "MIMDRAM: An End-to-End Processing-Using-DRAM System for High-Throughput, Energy-Efficient and Programmer-Transparent Multiple-Instruction Multiple-Data Computing," in *HPCA*, 2024.
- [38] M. M. Rehman, H. M. M. ur Rehman, W. Y. Kim, S. S. H. Sherazi, M. W. Rao, M. Khan, and Z. Muhammad, "Biomaterial-Based Nonvolatile Resistive Memory Devices Toward Ecofriendliness and Biocompatibility," *Appl. Electronic Mater.*, Jul. 2021.
- [39] A. P. Sample, D. J. Yeager, P. S. Powladge, A. V. Mamishev, and J. R. Smith, "Design of an RFID-Based Battery-Free Programmable Sensing Platform," *TIM*, Nov. 2008.
- [40] V. Seshadri *et al.*, "Ambit: In-Memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology," in *MICRO*, 2017.
- [41] V. Seshadri and O. Mutlu, "Simple Operations in Memory to Reduce Data Movement," in *Advances in Computers*. Elsevier, 2017, vol. 106.
- [42] A. Shafiee, A. Nag, N. Muralimanohar, R. Balasubramonian, J. P. Strachan, M. Hu, R. S. Williams, and V. Srikumar, "ISAAC: A Convolutional Neural Network Accelerator With In-Situ Analog Arithmetic in Crossbars," in *ISCA*, 2016.
- [43] S.-S. Sheu *et al.*, "Fast-Write Resistive RAM (RRAM) for Embedded Applications," *Design & Test of Computers*, Jan.-Feb. 2011.
- [44] P. Taylor, "Number of Wide-Area and Short-Range IoT Devices Worldwide From 2014 to 2031," <https://www.statista.com/statistics/1016276/wide-area-and-short-range-iot-device-installed-base-worldwide/>, 2025.
- [45] M. S. Q. Truong *et al.*, "The Memory Processing Unit: A Generalized Interface for End-to-End In-Memory Execution," in *HPCA*, 2026.
- [46] M. S. Q. Truong, E. Chen, D. Su, A. Glass, L. Shen, L. R. Carley, J. A. Bain, and S. Ghose, "RACER: Bit-Pipelined Processing Using Resistive Memory," in *MICRO*, 2021.
- [47] M. S. Q. Truong, L. Shen, A. Glass, A. Hoffmann, L. R. Carley, J. A. Bain, and S. Ghose, "Adapting the RACER Architecture to Integrate Improved In-ReRAM Logic Primitives," in *MICRO*, 2021.
- [48] United States Dept. of State, "Conflict Minerals," <https://www.state.gov/conflict-minerals>, 2026.
- [49] D. Vasisht, J. Shenoy, and R. Chandra, "L2D2: Low Latency Distributed Downlink for Low Earth Orbit Satellites," in *SIGCOMM*, 2021.
- [50] L. Wang, Y. Wang, W. Li, Z. Gao, Y. Han, and D. Wen, "Physically Transient Gelatin-Based Memristors of Buildable Logic Gates," *Gels*, Jun. 2025.
- [51] R. Wong, B. Feinberg, and S. Ghose, "DARTH-PUM: A Hybrid Processing-Using-Memory Architecture," in *ASPLOS*, 2026.
- [52] J. Xia, Z. Zhang, H. He, Y. Xu, D. Dong, R. Yang, and X. Miao, "Environment-Friendly Regenerated Cellulose Based Flexible Memristive Device," *Appl. Phys. Lett.*, Nov. 2021.
- [53] J. Yong *et al.*, "A Silk Fibroin Bio- Solution Processable Memristor," *Scientific Reports*, Nov. 2017.
- [54] J. X. Zhu, W. L. Zhou, Z. Q. Wang, H. Y. Xu, Y. Lin, W. Z. Liu, J. G. Maa, and Y. C. Liu, "Flexible, Transferable and Conformal Egg Albumen Based Resistive Switching Memory Devices," *RSC Advances*, Jun. 2017.