Global Research Journal in Engineering, Science, Technology and Management (GRJESTM)
Volume 02, Issue 01 (January-February) 2026, pp. 298-308
GRJESTM :: An odaswa™ Journal
ISSN : 3107-3611 (Online)

REVIEW ARTICLE

Designing Energy-Efficient Intelligent Systems Using Probabilistic Computing and p-Bits

Ananya Shetty*[1], Raghavendra Bhat[1], Karthick. M[2], and Nivedita Rao[2]

Abstract. The rapid expansion of artificial intelligence (AI), Internet of Things (IoT), edge computing, and autonomous systems has significantly increased the demand for energy-efficient computing technologies capable of supporting intelligent decision-making under stringent power constraints. Conventional deterministic computing architectures based on CMOS technology and the von Neumann model face growing challenges in terms of energy consumption, memory bottlenecks, and computational scalability. Probabilistic computing has emerged as a promising alternative paradigm by exploiting controlled stochastic behavior rather than deterministic binary operations. Central to this approach are probabilistic bits (p-bits), which fluctuate between binary states with tunable probabilities, enabling efficient implementation of optimization, inference, machine learning, and combinatorial computing tasks. Unlike quantum bits, p-bits operate at room temperature using existing semiconductor technologies, making them practical for near-term deployment. This review presents the fundamental principles of probabilistic computing and p-bit architectures, emphasizing their role in developing ultra-low-power intelligent systems. The paper discusses the theoretical foundations of probabilistic computation, stochastic information processing, p-bit devices, probabilistic circuits, and hybrid probabilistic-classical computing frameworks. By examining these foundational concepts, the review provides a comprehensive understanding of how probabilistic computing offers a scalable and energy-efficient alternative for future intelligent computing systems operating in resource-constrained environments.

Keywords: Probabilistic Computing, p-Bits, Stochastic Computing, Edge Artificial Intelligence, Ultra-Low-Power Computing, Intelligent Systems

1* School of Computer Science and Engineering, Manipal Academy of Higher Education, Manipal Institute of Technology Campus, Manipal, Karnataka, India
2 School of Engineering, Dayananda Sagar University, Devarakaggalahalli, Harohalli, Bengaluru, Karnataka, India

1. Introduction

Artificial Intelligence Artificial Intelligence (AI) has become one of the primary drivers of scientific innovation and digital transformation, enabling intelligent systems to perform complex tasks across healthcare, manufacturing, finance, transportation, robotics, and environmental monitoring. Recent advances in machine learning and deep learning have significantly improved the capability of computers to recognize patterns, make predictions, and support autonomous decision-making. However, these achievements have also resulted in rapidly increasing computational complexity, requiring massive processing power, high memory bandwidth, and substantial energy consumption during both training and inference [1], [2]. As AI applications continue to expand toward edge devices and embedded platforms, the demand for energy-efficient computing architectures has become increasingly important.

Most contemporary computing systems rely on deterministic binary logic implemented through Complementary Metal-Oxide-Semiconductor (CMOS) technology and the traditional von Neumann architecture, in which computation and memory are physically separated. This architectural organization requires continuous movement of data between processors and memory units, creating the well-known von Neumann bottleneck, which increases latency, energy consumption, and computational overhead [3]. Although hardware accelerators such as Graphics Processing Units (GPUs), Tensor Processing Units (TPUs), and neuromorphic processors have improved computational performance for AI workloads, the growing size and complexity of modern machine learning models continue to challenge existing computing infrastructures [4].

Emerging computing paradigms have therefore attracted significant attention as alternatives to conventional deterministic architectures. Among these, probabilistic computing has gained considerable interest because many real-world computational problems—including optimization, probabilistic inference, uncertainty modeling, and machine learning—are inherently stochastic rather than deterministic [5]. Instead of suppressing randomness as unwanted noise, probabilistic computing intentionally utilizes controlled stochastic behavior to perform computation. By representing uncertainty directly within hardware, probabilistic systems can solve complex optimization and inference problems with substantially lower energy consumption than conventional digital processors [6].

The fundamental computational element of probabilistic computing is the probabilistic bit (p-bit). Unlike a classical binary bit, which assumes either 0 or 1, or a quantum bit (qubit), which exists in coherent quantum superposition, a p-bit rapidly fluctuates between binary states according to a tunable probability distribution [7]. The probability of observing either state can be dynamically controlled through external inputs, allowing p-bits to represent uncertain information while remaining compatible with room-temperature semiconductor technologies. This characteristic distinguishes p-bits from qubits, which require quantum coherence and often operate under cryogenic conditions [8].

Networks of interconnected p-bits form probabilistic circuits (p-circuits) capable of solving computationally demanding problems through stochastic exploration of solution spaces. Rather than computing a single deterministic output, p-circuits continuously generate candidate solutions whose probabilities evolve according to network interactions and external constraints. Such behavior naturally aligns with probabilistic graphical models, Bayesian inference, Boltzmann machines, and combinatorial optimization algorithms, making p-bit architectures particularly attractive for machine learning and artificial intelligence applications [9].

One of the major advantages of probabilistic computing is its exceptional energy efficiency. Since p-bit devices rely on low-energy stochastic switching rather than deterministic logic transitions, they can perform many computational tasks with significantly reduced power consumption [10]. This capability is especially valuable for edge computing environments where battery capacity, thermal limitations, and hardware resources are constrained. Intelligent sensors, wearable healthcare devices, autonomous robots, Internet of Things (IoT) platforms, and environmental monitoring systems increasingly require computational architectures capable of supporting continuous learning and decision-making while operating under strict energy budgets [11].

Recent developments in spintronics, magnetic tunnel junctions (MTJs), resistive memory devices, and stochastic nanomagnetic technologies have accelerated practical implementation of p-bit hardware [12]. Low-barrier nanomagnets exhibit intrinsic thermal fluctuations that naturally generate probabilistic binary outputs, providing an efficient physical realization of p-bits without requiring complex random number generators. Furthermore, advances in CMOS-compatible stochastic circuits have demonstrated that probabilistic processors can be integrated with existing semiconductor technologies, enabling hybrid computing architectures that combine deterministic and probabilistic computation [13].

Understanding the fundamental principles of probabilistic computing is essential for appreciating its potential as a next-generation computing paradigm. Concepts such as stochastic computation, probabilistic inference, tunable random variables, p-bit architectures, probabilistic logic circuits, and hybrid probabilistic-classical computation provide the theoretical foundation for developing energy-efficient intelligent systems capable of solving complex computational problems under uncertainty. These principles also establish an important bridge between conventional digital computing, neuromorphic computing, and emerging unconventional computing technologies.

This paper presents a review of the fundamental principles of probabilistic computing and p-bit architectures with emphasis on their role in enabling ultra-low-power intelligent systems. It examines the theoretical concepts underlying stochastic information processing, probabilistic bits, probabilistic logic circuits, hardware implementations, and hybrid computing frameworks. By focusing on these foundational principles, the paper provides a comprehensive understanding of how probabilistic computing can contribute to the development of future intelligent systems that combine computational efficiency, adaptability, and energy sustainability.

2. Principles of Probabilistic Computing

Probabilistic computing is an emerging computational paradigm that intentionally incorporates randomness into information processing rather than treating it as an undesirable characteristic. Unlike conventional deterministic computing, where identical inputs always produce identical outputs, probabilistic computing allows computational states to evolve according to probability distributions. This stochastic behavior closely resembles many natural phenomena and provides an efficient framework for solving optimization, inference, machine learning, and uncertainty modeling problems [14]. As many real-world computational tasks involve incomplete information and probabilistic reasoning, probabilistic computing offers a promising alternative to traditional binary architectures by enabling intelligent systems to operate efficiently under uncertainty [15].

The foundation of probabilistic computing lies in representing computational variables as tunable random variables whose states continuously fluctuate according to external inputs and network interactions. Rather than producing a single deterministic output, probabilistic systems explore multiple candidate

solutions through stochastic state transitions, gradually converging toward the most probable solution. This computational strategy naturally aligns with Bayesian inference, Markov processes, Boltzmann statistics, and stochastic optimization techniques, making probabilistic computing particularly attractive for artificial intelligence and scientific computing applications [16].

2.1. Foundations of Probabilistic Computing

Traditional digital computers process information using deterministic binary logic, where each bit assumes a fixed value of either 0 or 1 until modified by logical operations. Although this approach has enabled remarkable advances in modern computing, it often becomes inefficient for problems involving uncertainty, optimization, and probabilistic reasoning. Many computational tasks in machine learning, pattern recognition, and scientific modeling require evaluating numerous possible solutions before identifying an optimal outcome, resulting in substantial computational complexity [17].

Probabilistic computing addresses these limitations by replacing deterministic binary variables with stochastic computational elements whose outputs fluctuate between logical states according to controllable probability distributions. Computation therefore becomes a statistical process rather than a purely deterministic sequence of logical operations. The probability associated with each computational state reflects the likelihood of a particular solution, allowing the system to naturally explore complex solution spaces while avoiding premature convergence to suboptimal solutions [7].

Unlike quantum computing, which relies on quantum coherence, entanglement, and superposition, probabilistic computing operates entirely within classical statistical mechanics and therefore does not require cryogenic environments or quantum error correction [18]. Instead, stochastic behavior is generated through thermal fluctuations, electronic noise, or intrinsically random physical devices. Consequently, probabilistic computing provides many computational advantages associated with stochastic optimization while remaining compatible with existing semiconductor manufacturing technologies.

Several characteristics distinguish probabilistic computing from conventional deterministic architectures:

- Controlled stochastic behavior.
- Probability-based information representation.
- Continuous exploration of multiple candidate solutions.
- Intrinsic support for uncertainty modeling.
- High computational efficiency for optimization and inference problems.

These properties establish probabilistic computing as a promising computing paradigm for future intelligent systems operating under limited computational resources.

2.2. Stochastic Information Processing

Stochastic information processing forms the theoretical basis of probabilistic computing. Rather than processing deterministic logical values, probabilistic systems manipulate probability distributions that evolve dynamically according to statistical interactions among computational elements [19]. This computational model closely resembles many natural processes in biological neural systems, statistical physics, and thermodynamic optimization.

One of the central concepts in stochastic information processing is the Boltzmann distribution, which describes the probability of observing various system states according to their corresponding energy levels. Computational systems utilizing stochastic dynamics tend to evolve toward lower-energy configurations while maintaining sufficient randomness to escape local minima during optimization [20]. This property enables probabilistic processors to efficiently solve combinatorial optimization problems that

are often computationally expensive for deterministic algorithms.

Another important principle is Markov Chain Monte Carlo (MCMC) sampling, which generates sequences of probabilistic states whose long-term distribution converges toward a desired probability distribution. MCMC techniques are widely employed in Bayesian inference, statistical learning, uncertainty quantification, and probabilistic graphical models [21]. Probabilistic hardware implementing stochastic state transitions can perform such sampling directly within hardware, reducing computational overhead associated with software-based simulations.

Probabilistic inference further extends stochastic information processing by estimating unknown variables based on observed evidence and prior probability distributions. Bayesian networks, factor graphs, and probabilistic graphical models provide mathematical frameworks for representing complex relationships among uncertain variables [22]. These models have found widespread applications in computer vision, speech recognition, medical diagnosis, autonomous systems, and intelligent decision support.

Compared with deterministic processors, stochastic information processing offers several important advantages:

- Efficient probabilistic inference.
- Improved robustness against hardware noise.
- Lower computational complexity for optimization tasks.
- Natural representation of uncertainty.
- Reduced energy consumption through approximate computation.

These advantages have positioned stochastic computing as an attractive computational model for future artificial intelligence systems requiring efficient probabilistic reasoning.

2.3.p-Bits and Their Operating Principles

The probabilistic bit (p-bit) is the fundamental computational element of probabilistic computing. A p-bit is a tunable stochastic binary device whose output continuously fluctuates between 0 and 1 according to an externally controllable probability distribution [7]. Unlike deterministic bits that remain fixed until explicitly modified, p-bits inherently generate random outputs while allowing the probability of each state to be adjusted through external bias signals.

Mathematically, the behavior of a p-bit is commonly represented using a sigmoidal activation function in which the output probability depends upon the weighted sum of external inputs [23]. As the input bias increases, the probability of observing one logical state gradually becomes dominant, while negative inputs favor the opposite state. When no external bias is present, both binary states occur with nearly equal probability due to intrinsic thermal fluctuations.

Physical implementations of p-bits have been demonstrated using several emerging device technologies, including:

- Low-barrier nanomagnets.
- Magnetic Tunnel Junctions (MTJs).
- Spintronic devices.
- CMOS stochastic circuits.
- Resistive switching devices.

Among these technologies, low-barrier nanomagnets have attracted considerable attention because thermal fluctuations naturally generate stochastic switching behavior while requiring extremely low operating energy [12]. When combined with CMOS control circuitry, these devices provide compact, scalable, and energy-efficient implementations of probabilistic processors [9].

Unlike qubits, which maintain coherent quantum states, p-bits operate entirely within classical physics and therefore avoid many of the technological challenges associated with quantum hardware. This practical compatibility with existing semiconductor fabrication processes

makes p-bit architectures attractive candidates for near-term deployment in intelligent embedded systems.

2.4. Probabilistic Logic Circuits

Individual p-bits become significantly more powerful when interconnected into probabilistic circuits (p-circuits). A p-circuit consists of multiple interacting p-bits whose stochastic outputs collectively evolve toward solutions satisfying predefined logical or optimization constraints [24]. Rather than implementing deterministic Boolean logic, p-circuits perform computation through statistical interactions among interconnected stochastic devices.

One remarkable property of p-circuits is invertible logic. Conventional digital logic gates compute outputs from known inputs in only one direction. In contrast, probabilistic circuits operate bidirectionally, allowing either inputs or outputs to be specified while automatically inferring unknown variables through probabilistic reasoning [7]. This capability greatly expands computational flexibility and enables efficient implementation of optimization, constraint satisfaction, and inverse problem solving.

Probabilistic logic circuits have successfully demonstrated implementations of:

- AND, OR, XOR, and NOT gates.
- Adders and multipliers.
- Integer factorization.
- Boolean satisfiability (SAT) problems.
- Traveling Salesman Problem (TSP).
- Ising model optimization.

Because computation emerges through stochastic exploration rather than exhaustive deterministic search, p-circuits often require substantially lower computational effort for solving combinatorial optimization problems [25]. Furthermore, their inherent fault tolerance allows reliable computation despite hardware variability and thermal noise, making probabilistic circuits particularly attractive for future ultra-low-power intelligent systems.

2.5. Hardware Technologies for p-Bit Architectures

The practical realization of probabilistic computing depends upon hardware devices capable of generating controllable stochastic behavior while consuming minimal energy. Unlike conventional digital circuits that suppress noise to achieve deterministic operation, p-bit architectures intentionally exploit intrinsic physical randomness to perform probabilistic computation. Recent advances in nanotechnology, spintronics, emerging memory devices, and CMOS-compatible stochastic circuits have significantly accelerated the development of practical p-bit hardware [26].

Among the various implementations, Low-Barrier Nanomagnets (LBNMs) have emerged as one of the most promising candidates for realizing p-bits. Unlike stable magnetic devices used in conventional memory applications, low-barrier nanomagnets possess reduced magnetic anisotropy, allowing thermal fluctuations at room temperature to continuously switch their magnetic orientation between two states [12]. These stochastic transitions naturally generate random binary outputs without requiring dedicated random number generators, making LBNMs highly energy-efficient computational elements.

Another important hardware technology is the Magnetic Tunnel Junction (MTJ), which consists of two ferromagnetic layers separated by an insulating tunnel barrier. The electrical resistance of an MTJ depends upon the relative orientation of the magnetic layers, enabling stochastic magnetic switching to be directly converted into probabilistic electrical outputs [9]. When combined with CMOS circuitry for input control and output sensing, embedded MTJs provide compact implementations of tunable p-bits suitable for large-scale probabilistic processors.

Beyond spintronic devices, several emerging memory technologies have also been investigated for probabilistic computing. These include:

- Resistive Random Access Memory (RRAM)
- Phase Change Memory (PCM)
- Ferroelectric Field-Effect Transistors (FeFETs)
- Spin-Orbit Torque (SOT) devices
- Memristive crossbar arrays

These technologies exhibit nonlinear switching dynamics, analog programmability, and low operating voltages that make them suitable for stochastic information processing [27]. In particular, memristive devices enable simultaneous storage and computation, thereby reducing data movement and mitigating the von Neumann bottleneck.

Although specialized nanodevices offer substantial energy savings, CMOS-based p-bit implementations remain attractive because they are compatible with existing semiconductor fabrication infrastructure. CMOS stochastic circuits emulate probabilistic behavior using compact random-number generators, feedback logic, and probabilistic activation functions [28]. This compatibility facilitates gradual integration of probabilistic computing into current digital computing platforms while supporting hybrid deterministic-probabilistic architectures.

As fabrication technologies continue to mature, large-scale arrays of interconnected p-bits are expected to provide scalable hardware platforms capable of supporting optimization, inference, machine learning, and intelligent decision-making with significantly reduced energy consumption.

2.6. Energy Efficiency and Ultra-Low-Power Intelligent Systems

Scientific discovery Energy efficiency represents one of the primary motivations for developing probabilistic computing architectures. Modern artificial intelligence systems require enormous computational resources for training and inference, resulting in increasing electricity consumption and thermal management challenges. Data centers supporting deep learning applications consume substantial amounts of electrical power, while edge devices often face severe limitations in battery capacity and processing capability [29].

Probabilistic computing addresses these challenges by replacing deterministic switching with stochastic computation that naturally exploits thermal fluctuations rather than overcoming them. Consequently, p-bit architectures can perform probabilistic inference and optimization while consuming considerably less energy than conventional deterministic processors [30].

Several architectural characteristics contribute to the energy efficiency of probabilistic systems:

- Low-voltage stochastic switching.
- Reduced computational complexity for optimization.
- Intrinsic parallel information processing.
- In-memory probabilistic computation.
- Approximate computation with graceful accuracy degradation.

Unlike deterministic processors that execute every logical operation precisely, probabilistic processors often compute approximate solutions sufficient for many artificial intelligence applications. This principle of approximate computing significantly reduces switching activity, computational latency, and energy consumption without substantially affecting overall system performance [31].

Another important advantage is the ability of probabilistic circuits to perform massively parallel stochastic exploration. Multiple p-bits simultaneously evaluate different candidate solutions, allowing optimization processes to converge efficiently without exhaustive sequential computation. This massively parallel architecture resembles biological neural systems, where large populations of neurons collectively process uncertain information using distributed stochastic dynamics [32].

The low-power characteristics of probabilistic computing make it particularly

attractive for numerous intelligent applications, including:

- Edge Artificial Intelligence (Edge AI)
- Internet of Things (IoT)
- Wearable healthcare devices
- Smart environmental monitoring
- Autonomous robotics
- Wireless sensor networks
- Smart agriculture
- Intelligent transportation systems

In these applications, battery lifetime, computational efficiency, and real-time decision-making are critical design requirements. Probabilistic processors provide an attractive balance between computational performance and energy efficiency, enabling intelligent systems to operate continuously under stringent resource constraints [11].

Furthermore, hybrid architectures combining deterministic processors with probabilistic accelerators allow computational tasks to be assigned according to their characteristics. Deterministic processors execute conventional arithmetic and control operations, whereas probabilistic hardware efficiently performs optimization, inference, sampling, and uncertainty estimation. Such heterogeneous architectures maximize overall computational efficiency while minimizing energy consumption [33].

2.7. Summary

Probabilistic computing represents a significant shift from conventional deterministic computation by embracing controlled randomness as a computational resource rather than treating it as undesirable noise. Through stochastic information processing, probabilistic inference, p-bits, and probabilistic logic circuits, this emerging paradigm provides an efficient framework for solving optimization, sampling, and machine learning problems that are computationally demanding for traditional digital architectures [7], [23].

Advances in spintronic devices, magnetic tunnel junctions, low-barrier nanomagnets, memristive technologies, and CMOS-compatible stochastic circuits have established practical hardware foundations for implementing scalable p-bit architectures [9], [12]. These devices exploit intrinsic thermal fluctuations to perform probabilistic computation while maintaining compatibility with existing semiconductor manufacturing processes.

One of the most compelling advantages of probabilistic computing is its exceptional energy efficiency. By combining approximate computation, massively parallel stochastic exploration, and hardware-native uncertainty representation, p-bit architectures significantly reduce computational complexity and power consumption. These characteristics position probabilistic computing as a promising technology for ultra-low-power intelligent systems operating in resource-constrained environments, including edge AI, autonomous robotics, IoT, wearable electronics, and scientific optimization.

As hardware technologies, probabilistic algorithms, and hybrid deterministic-probabilistic computing frameworks continue to evolve, p-bit architectures are expected to become increasingly important components of next-generation intelligent computing systems. Their ability to integrate efficient probabilistic reasoning with scalable semiconductor technologies offers a compelling pathway toward sustainable, adaptive, and energy-efficient artificial intelligence.

3. Conclusion

Probabilistic computing has emerged as a promising post-CMOS computing paradigm that fundamentally redefines the role of randomness in computation. Unlike conventional deterministic architectures that attempt to eliminate stochastic behavior, probabilistic computing intentionally exploits controlled randomness as a computational resource for solving optimization, inference, and machine learning problems. By utilizing probabilistic bits (p-bits) as tunable

stochastic computational elements, this emerging paradigm offers an energy-efficient alternative to conventional digital computing while maintaining compatibility with existing semiconductor technologies.

The principles discussed in this review—including stochastic information processing, probabilistic inference, p-bit architectures, probabilistic logic circuits, and hybrid deterministic-probabilistic computing—collectively establish the theoretical foundation of probabilistic computing. These concepts demonstrate that randomness can significantly improve computational efficiency for numerous artificial intelligence tasks, particularly those involving uncertainty, probabilistic reasoning, combinatorial optimization, and approximate computation. Unlike quantum computing, which requires sophisticated quantum hardware and cryogenic operating environments, p-bit architectures operate at room temperature using CMOS-compatible spintronic and nanomagnetic devices, making them attractive candidates for near-term practical deployment.

Recent advances in low-barrier nanomagnets, magnetic tunnel junctions, memristive devices, stochastic CMOS circuits, and probabilistic spin logic have accelerated the development of scalable hardware platforms capable of implementing large networks of interconnected p-bits. These technologies enable massively parallel stochastic computation while significantly reducing energy consumption compared with deterministic processors. Furthermore, hybrid computing frameworks that integrate conventional digital processors with probabilistic accelerators provide a practical pathway toward deploying probabilistic computing in existing computing infrastructures.

The application potential of probabilistic computing extends across numerous scientific and engineering domains. Optimization problems in logistics, manufacturing, finance, communication networks, and scientific research can benefit from probabilistic search mechanisms that efficiently explore large solution spaces. Similarly, artificial intelligence applications—including Bayesian learning, probabilistic graphical models, neural networks, autonomous robotics, computer vision, natural language processing, and edge AI—can exploit p-bit architectures for low-power intelligent inference under uncertain conditions. These characteristics make probabilistic computing particularly attractive for Internet of Things (IoT) devices, wearable electronics, autonomous vehicles, smart healthcare systems, environmental monitoring, and intelligent sensor networks where computational resources and battery capacity are limited.

Despite these promising developments, several research challenges remain before large-scale probabilistic computing systems become commercially viable. Device variability, stochastic stability, fabrication consistency, hardware scalability, programming methodologies, and software-hardware co-design require further investigation. Developing standardized programming frameworks, compiler support, benchmarking methodologies, and application-specific probabilistic algorithms will be essential for accelerating industrial adoption. Continued research in spintronic materials, nanoscale device engineering, probabilistic architectures, and machine learning algorithms is expected to address many of these limitations over the coming years.

In conclusion, probabilistic computing represents an important milestone in the evolution of next-generation computing technologies. By combining stochastic hardware, probabilistic reasoning, and energy-efficient computation, p-bit architectures provide a practical and scalable foundation for developing ultra-low-power intelligent systems. As research progresses, probabilistic computing is expected to complement conventional digital processors, neuromorphic systems, and quantum computing platforms, thereby contributing significantly to the future landscape of sustainable artificial intelligence and intelligent computing.

Acknowledgement

The authors sincerely acknowledge the guidance and support of the faculty members of their institution for their valuable insights and encouragement throughout this study. Special thanks are extended to the library and research staff for providing access to relevant resources. The authors also express gratitude to colleagues and peers for their constructive feedback and motivation during the completion of this work.

References

[1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553, pp. 436–444, 2015.

[2] E. Brynjolfsson and A. McAfee, "The business of artificial intelligence," MIT Sloan Management Review, vol. 59, no. 2, pp. 1–14, 2017.

[3] J. Backus, "Can programming be liberated from the von Neumann style? A functional style and its algebra of programs," Communications of the ACM, vol. 21, no. 8, pp. 613–641, 1978.

[4] N. P. Jouppi et al., "In-datacenter performance analysis of a Tensor Processing Unit," in Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA), 2017, pp. 1–12.

[5] S. C. Goldstein, J. D. Campbell, and T. C. Mowry, "Probabilistic computation," Computer, vol. 38, no. 8, pp. 24–30, 2005.

[6] P. Date, "Probabilistic computing: Hype or reality?" Communications of the ACM, vol. 63, no. 7, pp. 32–34, 2020.

[7] K. Y. Camsari, R. Faria, B. M. Sutton, and S. Datta, "Stochastic p-bits for invertible logic," Physical Review X, vol. 7, no. 3, Art. no. 031014, 2017.

[8] J. Preskill, "Quantum computing in the NISQ era and beyond," Quantum, vol. 2, p. 79, 2018.

[9] K. Y. Camsari, S. Salahuddin, and S. Datta, "Implementing p-bits with embedded MTJ," IEEE Electron Device Letters, vol. 38, no. 12, pp. 1767–1770, 2017.

[10] S. Datta, "Probabilistic spin logic," Applied Physics Letters, vol. 119, no. 18, Art. no. 180501, 2021.

[11] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637–646, 2016.

[12] K. Y. Camsari, B. M. Sutton, and S. Datta, "Low-barrier nanomagnets as hardware p-bits for probabilistic spin logic," Applied Physics Reviews, vol. 6, no. 1, Art. no. 011305, 2019.

[13] R. Faria, K. Y. Camsari, and S. Datta, "Hardware implementation of probabilistic circuits using stochastic p-bits," IEEE Journal on Exploratory Solid-State Computational Devices and Circuits, vol. 4, no. 2, pp. 52–60, 2018.

[14] J. von Neumann, "Probabilistic logics and the synthesis of reliable organisms from unreliable components," Automata Studies, Princeton University Press, pp. 43–98, 1956.

[15] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Hoboken, NJ, USA: Pearson, 2021.

[16] D. J. C. MacKay, Information Theory, Inference, and Learning Algorithms. Cambridge, UK: Cambridge University Press, 2003.

[17] C. H. Bennett, "Logical reversibility of computation," IBM Journal of Research and Development, vol. 17, no. 6, pp. 525–532, 1973.

[18] M. A. Nielsen and I. L. Chuang, Quantum Computation and Quantum Information. Cambridge, UK: Cambridge University Press, 2010.

[19] H. C. Berg, Random Walks in Biology. Princeton, NJ, USA: Princeton University Press, 1993.

[20] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi, "Optimization by simulated annealing," Science, vol. 220, no. 4598, pp. 671–680, 1983.

[21] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," Biometrika, vol. 57, no. 1, pp. 97–109, 1970.

[22] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. San Mateo, CA, USA: Morgan Kaufmann, 1988.

[23] K. Y. Camsari, B. M. Sutton, R. Faria, and S. Datta, "p-bits for probabilistic spin logic," Applied Physics Reviews, vol. 6, no. 1, Art. no. 011305, 2019.

[24] K. Y. Camsari, S. Salahuddin, and S. Datta, "Implementing p-circuits with stochastic p-bits," IEEE Journal on Exploratory Solid-State Computational Devices and Circuits, vol. 5, no. 1, pp. 15–23, 2019.

[25] S. Datta, K. Y. Camsari, and B. M. Sutton, "Probabilistic spin logic via low-barrier

nanomagnets," Scientific Reports, vol. 9, Article 15599, 2019.

[26] S. Datta and K. Y. Camsari, "Hardware probabilistic computing with p-bits," Nature Electronics, vol. 5, no. 10, pp. 604–615, 2022.

[27] G. W. Burr et al., "Neuromorphic computing using non-volatile memory," Advances in Physics: X, vol. 2, no. 1, pp. 89–124, 2017.

[28] K. Y. Camsari, R. Faria, S. Salahuddin, and S. Datta, "Stochastic computing with p-bits," IEEE Design & Test, vol. 36, no. 3, pp. 53–61, 2019.

[29] D. Patterson et al., "Carbon emissions and large neural network training," arXiv:2104.10350, 2021.

[30] V. De, "Energy-efficient computing: From devices to architectures," Proceedings of the IEEE, vol. 109, no. 5, pp. 669–688, 2021.

[31] S. Venkataramani, A. Ranjan, K. Roy, and A. Raghunathan, "Approximate computing and the quest for computing efficiency," Proceedings of the 52nd Design Automation Conference, 2015.

[32] W. Maass, "Networks of spiking neurons: The third generation of neural network models," Neural Networks, vol. 10, no. 9, pp. 1659–1671, 1997.

[33] K. Y. Camsari, S. Salahuddin, and S. Datta, "p-Bits for probabilistic computing: Theory and applications," Proceedings of the IEEE, vol. 109, no. 5, pp. 761–781, 2021.

[34] S. Datta, Lessons from Nanoelectronics: A New Perspective on Transport. Singapore: World Scientific, 2012.

[35] J. M. Martinis and M. Nahum, "Probabilistic computing with stochastic nanomagnets," Nature Nanotechnology, vol. 15, no. 7, pp. 507–508, 2020.

[36] K. Y. Camsari, B. M. Sutton, and S. Datta, "Probabilistic computing with p-bits," Physics Today, vol. 72, no. 8, pp. 38–43, 2019.

[37] G. E. Hinton and T. J. Sejnowski, "Learning and relearning in Boltzmann machines," in Parallel Distributed Processing: Explorations in the Microstructure of Cognition, vol. 1, D. E. Rumelhart and J. L. McClelland, Eds. Cambridge, MA, USA: MIT Press, 1986, pp. 282–317.

[38] J. S. Denker, "Neural networks and annealing methods," Complex Systems, vol. 1, no. 6, pp. 877–922, 1987.

[39] D. J. Amit, Modeling Brain Function: The World of Attractor Neural Networks. Cambridge, UK: Cambridge University Press, 1989.

[40] T. N. Theis and H.-S. P. Wong, "The end of Moore's law: A new beginning for information technology," Computing in Science & Engineering, vol. 19, no. 2, pp. 41–50, 2017.

[41] International Energy Agency (IEA), Digitalisation and Energy. Paris, France: IEA Publications, 2017.

[42] M. Horowitz, "Computing's energy problem (and what we can do about it)," in IEEE International Solid-State Circuits Conference (ISSCC), 2014, pp. 10–14.