Qwen

Beyond Brute Force: A Synthesis of Guided Search Algorithms and Hardware Acceleration for Combinatorial Problems

The challenge of navigating extremely large state spaces, often referred to as combinatorial or state space explosion, represents a fundamental barrier to solving many complex problems across science and engineering [92] [111]. When the number of possible configurations or solutions grows exponentially with the problem size, exhaustive enumeration—brute force—is computationally intractable. This report provides a comprehensive analysis of non-brute-force methods designed to find a single optimal solution within these vast domains. It explores a spectrum of algorithmic strategies, from classical informed search and metaheuristics to advanced mathematical transformations and learning-based guidance. Concurrently, it examines the specialized hardware architectures developed to accelerate these guided searches, including theoretical models, emerging photonic and neuromorphic systems, and established custom accelerators. The research focuses on application domains such as artificial intelligence, cryptographic key recovery, and prime number discovery, while explicitly excluding quantum computing. The central theme is the synergistic relationship between intelligent algorithms and specialized hardware, which together make previously intractable search problems computationally feasible.

Algorithmic Strategies for Guided Search in Vast State Spaces

To overcome the prohibitive computational cost of brute-force search, a diverse array of algorithmic strategies has been developed. These methods aim to intelligently navigate the state space, either by reducing its effective size or by guiding the search toward regions that are more likely to contain the optimal solution. These strategies range from foundational heuristic search techniques to sophisticated machine learning-driven approaches and profound mathematical reformulations of the underlying problems. Their application is critical in domains like artificial intelligence planning, where an agent must find a sequence of actions to achieve a goal; in cryptanalysis, where an attacker seeks a

secret key; and in number theory, where the task is to identify a prime number among a vast set of candidates [6] [24] [41].

Classical informed search algorithms represent a foundational layer of guided search. Unlike uninformed methods that explore the state space without any domain-specific knowledge, informed algorithms utilize a heuristic function to estimate the cost or distance from a given state to the goal [79]. This guidance significantly improves search efficiency. The most prominent of these is the A *search algorithm, which guarantees finding an optimal solution if the heuristic function is admissible, meaning it never overestimates the true cost to reach the goal* [80] [81]. *A* evaluates nodes based on the sum of two components: the actual cost from the start node to the current node, denoted $g(n)$, and the heuristic estimate of the cost from the current node to the goal, denoted $h(n)$ [79]. Another variant, Greedy Best-First Search, uses only the heuristic component $h(n)$, prioritizing paths that appear to be closest to the goal at each step. While this can lead to very fast solutions, it does not guarantee optimality and can easily get trapped in local optima by always choosing the most immediately promising path [79]. For problems with enormous state spaces where storing all explored nodes is impractical, Iterative Deepening A *(IDA)* offers a compelling alternative. It combines the optimality and completeness guarantees of A* with the low memory usage of depth-first search by performing a series of depth-limited searches with progressively increasing cost limits defined by the $f(n)=g(n)+h(n)$ function [79]. In scenarios where even exploring the best few paths at each level is too computationally expensive, Beam Search provides a practical compromise. It functions as a constrained version of Best-First Search, retaining only a fixed number of the most promising nodes (the 'beam width') at each level of the search, thus trading off completeness and optimality for significant gains in efficiency. This makes it particularly suitable for applications like machine translation, where computational constraints necessitate a balance between search depth and resource usage [79] [92].

When finding an exact optimal solution is deemed too difficult or time-consuming, metaheuristic algorithms provide powerful frameworks for finding high-quality approximate solutions. These are high-level strategies that operate on an iterative improvement mechanism. Simulated Annealing is a well-known example, inspired by the physical process of annealing in metallurgy [103]. It is an iterative probabilistic technique that allows for occasional "uphill" moves that worsen the current solution, thereby enabling the search to escape local optima and explore a broader portion of the search space [103]. Other prominent metaheuristics include Evolutionary Algorithms, which mimic the principles of natural selection and genetics, evolving a population of candidate solutions through operations like mutation and crossover over successive generations [109].

These algorithms are highly versatile and have been applied to train neural networks and solve complex optimization problems [109]. The choice of algorithm depends heavily on the problem's structure and the required trade-off between solution quality and computational resources.

A more profound strategy involves transforming the original problem into a different, mathematically equivalent but potentially more tractable form. A recurring and powerful technique in this category is mapping combinatorial optimization problems onto the Ising model or its equivalent, Quadratic Unconstrained Binary Optimization (QUBO). Many NP-hard problems, such as Max-Cut, can be converted into an Ising problem where the goal is to find the ground state—the configuration of interacting spins with the lowest possible energy [25] [28] [112]. Solving this physical system becomes synonymous with solving the original computational problem [31]. This transformation is a unifying principle that enables a wide range of novel hardware solvers, from photonic oscillators to neuromorphic processors, to be applied to diverse optimization tasks [32]. Before any search is initiated, another crucial pre-processing step is symmetry breaking. Many problems possess symmetries where multiple distinct states correspond to the same solution. By adding explicit constraints to the problem formulation, some of these symmetric solutions can be made unacceptable, drastically reducing the search space without eliminating any optimal solutions [2]. Research has led to certification methods based on proof systems that can generate machine-verifiable certificates to ensure the correctness of solutions computed using these techniques, enhancing their reliability [13].

In recent years, machine learning (ML) and deep learning (DL) have emerged as potent tools for guiding search and analyzing complex systems, particularly in the field of cryptanalysis. A notable advancement is the development of "neural distinguishers." In traditional cryptanalysis, an attacker might use statistical properties to differentiate the output of a cipher from random data. A neural distinguisher replaces this manual process with a trained deep neural network, which can learn far more subtle patterns [7] [91]. Pioneered by Gohr in CRYPTO 2019, this approach demonstrated superior performance in key-recovery attacks against ciphers like the NSA's Speck [35] [37] [40]. The neural network acts as a highly accurate distinguisher, allowing the attacker to focus their efforts on the most promising parts of the search space. Modern, state-of-the-art attacks are increasingly hybrid, combining classical cryptanalytic techniques, such as differential-linear cryptanalysis, with ML. The neural distinguisher refines the attack, leading to significant improvements in key recovery rates compared to purely classical methods [8] [9] [38] [39]. This fusion of classical algorithmic insight with the pattern-recognition power of deep learning represents a new frontier in tackling large search spaces in cryptography. Beyond security, this learning-based approach extends to AI itself. The development of

Spiking Neural Networks (SNNs), which emulate the brain's event-driven communication, is being paired with neuromorphic hardware through a process known as algorithm-hardware co-design. This involves creating SNNs that are specifically tailored to exploit the low-power, asynchronous nature of neuromorphic processors, making the search for optimal network parameters or model configurations vastly more efficient [48] [50] [51].

Theoretical and Realized Architectures for Combinatorial Traversal

Addressing the immense challenge of combinatorial search requires not only intelligent algorithms but also specialized hardware capable of executing them at scale. The limitations of traditional von Neumann architectures, where processing and memory are separate, create a bottleneck that hinders the performance of massively parallel search algorithms. Consequently, research has focused on developing alternative computational paradigms that directly map the problem's structure onto the physical hardware, aiming for unprecedented levels of parallelism and efficiency. The ideal theoretical architecture envisioned for this purpose is a programmable analog Ising machine, while several real-world physical platforms—including photonic, neuromorphic, and memristive systems—are actively being developed to realize its principles.

The concept of an ideal theoretical architecture is rooted in analog computing, specifically the idea of an Ising machine. The core principle is to circumvent the bottlenecks inherent in digital von Neumann computing by creating a physical system whose natural dynamics evolve toward the solution of a given combinatorial optimization problem [26] [29]. In this model, a problem is formulated as an Ising or QUBO model, and the hardware is engineered so that its ground state (lowest energy configuration) corresponds directly to the optimal solution [25]. An ideal machine would possess several key features. First, it must offer full programmability, allowing it to solve arbitrary Ising problems, not just specific instances like Max-Cut [104]. Second, it requires all-to-all connectivity, meaning every variable (spin) in the problem can interact with every other variable, a feature essential for representing general graphs [83]. Third, scalability is paramount, with the ability to support a large number of variables, with research demonstrating capabilities up to 20,000-node problems [83]. Finally, it must operate with high parallelism, simultaneously updating all variables to rapidly converge towards the solution.

Several physical platforms are being explored to build these machines. Photonic Ising Machines (CIMs) are a leading approach, utilizing networks of coupled optical oscillators to simulate the Ising model [23] [115]. These systems leverage the speed of light and the inherent parallelism of optics to perform computations at high speeds [46]. Different implementations include coherent Ising machines (CIMs) and spatial-photonic Ising machines (SPIMs) [24] [59]. Recent advancements, such as the use of focal-plane division, have enabled the creation of fully programmable SPIMs capable of solving arbitrary Ising problems, overcoming earlier limitations related to connectivity and programmability [104]. Despite their promise, challenges remain in areas like precision and control [114].

Neuromorphic computing offers another pathway, focusing on silicon-based electronic systems designed to mimic the structure and function of biological brains [20] [89]. These systems are characterized by high parallelism, low power consumption, and event-driven computation, making them well-suited for running algorithms like Spiking Neural Networks (SNNs) [50] [51]. Specific neuromorphic architectures have been developed to solve Ising problems. For instance, CMOS annealing uses standard CMOS technology to perform an annealing-like process to find optimal solutions [11]. A more advanced example is NeuroSA, a neuromorphic architecture designed to ensure asymptotic convergence to the Ising ground state [44] [61]. It leverages asynchronous hardware and techniques like Fowler-Nordheim tunneling for annealing, combined with autoencoders to define the solution space, representing a sophisticated implementation of the Ising machine concept in silicon [60] [62] [84].

Memristive systems represent a third major hardware paradigm. Memristors are non-volatile memory devices whose resistance can be programmed to store information [88]. They are being integrated into analog in-memory computing architectures, where computation occurs directly within the memory array, eliminating the need to shuttle data back and forth to a processor [18] [67]. These systems can implement dynamical systems like Hopfield networks to solve optimization problems [113]. Research indicates that memristor-based hardware can outperform quadratic Ising machines, offering potential for implementing more complex, higher-order Ising models [63]. Such co-design and optimization have been the main focus of memristive neuromorphic engineering [18]. Alongside these emerging paradigms, established technologies like Application-Specific Integrated Circuits (ASICs) and Field-Programmable Gate Arrays (FPGAs) continue to play a vital role in accelerating search tasks. ASICs are custom-designed for specific applications, offering maximum performance and efficiency [15]. For example, ASICs are being developed to accelerate post-quantum cryptographic algorithms like those based on lattices, providing order-of-magnitude speedups for critical operations [72] [73] [102]. FPGAs

serve as a flexible platform for prototyping and deploying custom logic, used for accelerating machine learning models, real-time cryptographic processing, and implementing various neuromorphic architectures [49] [74] [106] [118]. Together, these diverse hardware platforms—from theoretical analog machines to practical ASICs—form a multi-pronged effort to build the computational engines needed to conquer combinatorial complexity.

Hardware Architecture Type	Core Principle	Example Implementations	Key Advantages	Key Challenges
Ideal Theoretical Machine	Direct physical mapping of problem to hardware dynamics to find the ground state.	Programmable Analog Ising Machine [26] [29]	All-to-all connectivity, massive parallelism, potential for exponential speedup.	Full programmability, scalability, and stability are still theoretical goals.
Photonic Ising Machines	Using networks of coupled optical oscillators to simulate spin interactions.	Coherent Ising Machines (CIMs), Spatial-Photonic Ising Machines (SPIMs) [23] [24]	High speed (light propagation), massive parallelism, low latency.	Precision, programmability for arbitrary problems, and control stability [46] [59] [114].
Neuromorphic Ising Machines	Silicon-based electronic systems mimicking brain structure for event-driven, low-power computation.	NeuroSA (using Fowler-Nordheim tunneling), CMOS Annealing [11] [44] [61]	High energy efficiency, asynchronous operation, scalable in silicon [89].	Ensuring asymptotic convergence, device variability, and integration complexity.
Memristive Systems	Using memristive devices in in-memory computing arrays to implement dynamical optimization solvers.	Memristive Hopfield Networks, Higher-order Ising Solvers [18] [63]	Eliminates von Neumann bottleneck, potential for ultra-low power [67], supports higher-order interactions.	Long-term stability, weight drift, and manufacturing variability [18].
ASICs / FPGAs	Custom or reconfigurable digital circuits optimized for specific algorithms.	Lattice-based crypto accelerators [15] [72], FPGA-based SNN accelerators [106] [108]	High performance, low latency, and energy efficiency for targeted tasks.	Lack of general-purpose flexibility (ASICs), limited logic resources (FPGAs).

Application Domains: From Cryptographic Keys to Prime Discovery

The theoretical underpinnings of guided search algorithms and specialized hardware find concrete expression in several critical application domains. The user-specified areas of cryptography, prime number discovery, and artificial intelligence all grapple with vast state spaces where the existence of a single correct answer—a cryptographic key, a large prime, or an optimal model configuration—must be found efficiently. In each domain,

the interplay between algorithmic innovation and hardware acceleration is reshaping what is computationally possible.

In cryptography, the task of cracking a key is a quintessential search problem defined by a massive state space. Historically, the security of public-key systems like RSA relies on the presumed difficulty of integer factorization [71]. The General Number Field Sieve (GNFS) is the most efficient classical algorithm known for this task [14]. Its performance is critically dependent on the initial stage of polynomial selection, a sub-problem that researchers continuously work to improve [3] [5]. Enhancements to the size and root properties of the polynomials used can lead to significant reductions in overall computation time [3] [4]. An alternative powerful method is the Elliptic Curve Method (ECM), which has seen record-setting performance when accelerated on graphics cards, demonstrating the importance of hardware in this domain [116]. The modern landscape of cryptanalysis has evolved beyond purely mathematical approaches. A significant breakthrough has been the integration of machine learning into classical cryptanalytic techniques. Pioneered by Gohr in 2019, ML-aided differential cryptanalysis uses a deep neural network as a highly accurate "distinguisher" [35] [37]. This neural distinguisher surpasses traditional statistical methods, allowing attackers to more effectively identify the internal state of a cipher and recover the secret key [91]. Hybrid attacks, which combine classical methods like differential-linear cryptanalysis with an ML-powered distinguisher, have shown substantial improvements in key recovery rates for round-reduced block ciphers [8] [9] [38]. This arms race between attack and defense drives the development of post-quantum cryptographic standards, for which hardware accelerators are already being designed to provide the necessary performance for secure cloud computing in a future where quantum computers may render current encryption obsolete [71] [73] [75].

Prime number discovery presents another classic problem of searching for a needle in a haystack. While primality testing can determine if a given number is prime, discovering new large primes is a more active search process. Traditional primality proving algorithms, such as the Elliptic Curve Primality Proving (ECPP) algorithm, offer a deterministic method to prove the primality of a large integer [41] [78]. A key feature of ECPP is that it produces a certificate of primality that can be independently and efficiently verified, providing a strong guarantee of correctness [42] [77]. This contrasts with probabilistic tests that can only assert primality with a very high degree of certainty. While ECPP is a powerful tool, some research explores more guided, heuristic-based frameworks for searching for primes. One paper introduces a transcendental geometric framework that aims to guide the search through resonance structures induced by logarithmic relationships, suggesting a move towards leveraging mathematical structure

to direct the search rather than relying solely on trial-and-error methods [43]. The search for extremely large primes, such as Mersenne primes, is often facilitated by distributed computing projects, but the underlying mathematical and algorithmic principles remain central to identifying these rare numbers.

In the realm of artificial intelligence and latent space exploration, the state space is typically the immense parameter space of a neural network or the embedding space of a generative model. Finding an optimal solution here could mean training a model to minimize a loss function or generating a specific type of output from a latent code. Classical search and optimization techniques are fundamental to training ANNs and DLs, and metaheuristics like genetic algorithms are frequently used for this purpose [109]. However, the scale of modern AI models makes these processes incredibly resource-intensive. To address this, there is a growing trend toward using Graph Neural Networks (GNNs) for combinatorial optimization. By treating a problem instance as a graph, a GNN can be trained end-to-end to produce a one-shot solution, effectively bypassing the need for iterative search altogether [107]. This represents a paradigm shift from dynamic search to learned prediction. This trend is deeply connected to the development of neuromorphic hardware. The push for more efficient AI has driven the design of Spiking Neural Networks (SNNs), which communicate via discrete events and are naturally suited for low-power, asynchronous hardware [50] [51]. The development of ACE-SNN, a co-designed SNN generated from a standard CNN, exemplifies this approach, where the algorithm is specifically crafted to run efficiently on neuromorphic processors [48]. Similarly, the development of frameworks like NeuroBench aims to create standardized benchmarks for evaluating neuromorphic algorithms and systems, fostering progress in this area [86]. These efforts illustrate a holistic approach where the entire pipeline, from algorithm design to hardware implementation, is optimized to make searching the vast space of AI model configurations more efficient and sustainable [68].

Co-Design and Performance Benchmarking of Next-Generation Search Engines

The development of next-generation systems for navigating large state spaces is increasingly defined by the principle of algorithm-hardware co-design. This paradigm moves beyond the traditional separation of software algorithms and hardware platforms, instead treating them as interconnected components of a single system to be optimized together [67] [70]. This approach is critical because the unique strengths and constraints of

emerging hardware architectures—such as neuromorphic processors or photonic annealers—can only be fully exploited when the algorithms running on them are specifically tailored to their operational characteristics. Concurrently, the proliferation of these novel platforms necessitates rigorous performance benchmarking to objectively evaluate their capabilities and compare them against both classical computing systems and each other.

Co-design is particularly evident in the field of neuromorphic computing. Here, the hardware is built around principles of biological neural systems, featuring massive parallelism, event-driven computation, and low power consumption [68] [89]. To harness these features, algorithms must also be brain-inspired. Spiking Neural Networks (SNNs) are the natural choice, as their communication mechanism of sending discrete "spikes" aligns perfectly with the asynchronous nature of neuromorphic chips [51]. The co-design process involves not just running existing algorithms on new hardware, but actively shaping the algorithm during its creation. For example, researchers develop SNNs from pre-trained conventional CNNs, using techniques like quantization-aware gradient descent to adapt the model for the spiking format before deployment [48]. This ensures that the resulting SNN is not only functionally similar to the original model but is also optimized for the specific computational primitives of the target neuromorphic substrate [66]. This tight coupling between the algorithm's structure and the hardware's capabilities is essential for achieving the promised gains in energy efficiency and performance, especially for applications in resource-constrained environments like robotics and IoT [17] [90].

This co-design philosophy extends beyond neuromorphism to other emerging hardware platforms. In the context of photonic Ising machines, the interaction weights of the Ising problem must be encoded into the optical system. This requires developing heuristic algorithms that can translate an abstract problem into the physical parameters of the machine, such as phase shifts and couplings, in a way that the system can solve it efficiently [82]. Similarly, in memristive in-memory computing, the analog nature of the devices imposes constraints on the numerical precision and stability of the computations. Therefore, algorithms must be designed to be robust to device non-idealities, and hardware-software co-design strategies are emerging to track and compensate for issues like weight drift caused by environmental changes [18] [19]. Even in more established domains like FPGA acceleration, co-design is a key practice. Designers can configure various hardware "knobs" in an HLS (High-Level Synthesis) tool to optimize an accelerator for a specific workload, balancing factors like throughput, latency, and power consumption [74].

As these diverse and specialized hardware platforms emerge, establishing objective performance metrics becomes crucial for assessing their true value. A simple measure like raw speed is insufficient; a comprehensive evaluation must consider accuracy, energy efficiency, scalability, and problem-specific metrics. To address this need, dedicated benchmarking frameworks are being developed. NeuroBench, for instance, is a community-driven framework designed to benchmark neuromorphic algorithms and systems, providing standardized testbeds for fair comparison [86]. In the context of Ising machines, researchers have developed universal, reproducible generative models for QUBO/Ising instances to benchmark and compare the performance of different hardware platforms [45]. These benchmarks allow for a quantitative assessment of how well a given machine solves problems of varying size and structure. For example, Fujitsu's Digital Annealer has been benchmarked on the Max-Cut problem, a canonical NP-hard problem that captures the main complexity of the broader QUBO class [112]. Comparing the performance of a coherent Ising machine against a D-Wave quantum annealer has also provided insights into the relative importance of factors like connectivity in these specialized systems [114]. Such systematic evaluations are essential for guiding future research and development, helping to identify the strengths and weaknesses of each architectural approach and informing decisions about where to invest resources. Ultimately, the synergy between co-designed algorithms and rigorous benchmarking will be the engine driving progress in building practical, powerful computers for tackling the world's most challenging search problems.

Reference

1. A Kerr soliton Ising machine for combinatorial optimization problems https://arxiv.org/html/2508.00810v1
2. [PDF] Symmetry in Constraint Programming - ResearchGate https://www.researchgate.net/profile/Jean-Francois-Puget/publication/240829678_Symmetry_in_Constraint_Programming/links/00b7d53b6948166dbe000000/Symmetry-in-Constraint-Programming.pdf
3. Better polynomials for GNFS - HAL-Inria https://inria.hal.science/hal-01089507v1
4. Improvements to the general number field sieve for discrete ... https://www.researchgate.net/publication/220576503_Improvements_to_the_general_number_field_sieve_for_discrete_logarithms_in_prime_fields_A_comparison_with_the_Gaussian_integer_method

5. Better polynomials for GNFS - ResearchGate https://www.researchgate.net/publication/312436132_Better_polynomials_for_GNFS
6. [PDF] Advancements and Prospects in Large Integer Factorization https://pdfs.semanticscholar.org/7003/0a159c209a34bffd9118dc5a5ce625407991.pdf
7. (PDF) Machine Learning and Cryptanalysis: An In-Depth Exploration ... https://www.researchgate.net/publication/378349918_Machine_Learning_and_Cryptanalysis_An_In-Depth_Exploration_of_Current_Practices_and_Future_Potential
8. Mixture Differential Cryptanalysis on Round-Reduced SIMON32/64 ... https://www.mdpi.com/2227-7390/12/9/1401
9. Improving deep learning-based neural distinguisher with multiple ... https://pmc.ncbi.nlm.nih.gov/articles/PMC12012056/
10. Recent Trend of Neuromorphic Computing Hardware: Intel's ... https://www.researchgate.net/publication/348979389_Recent_Trend_of_Neuromorphic_Computing_Hardware_Intel's_Neuromorphic_System_Perspective
11. Higher-order neuromorphic Ising machines—autoencoders and ... https://www.nature.com/articles/s41467-026-71937-4
12. Photonic Ising machines for combinatorial optimization problems https://www.researchgate.net/publication/384911822_Photonic_Ising_machines_for_combinatorial_optimization_problems
13. Certified Symmetry and Dominance Breaking for Combinatorial ... https://arxiv.org/abs/2203.12275
14. On General Number Field Sieve and its Polynomial Selection https://www.researchgate.net/publication/389977247_On_General_Number_Field_Sieve_and_its_Polynomial_Selection
15. Lattice-Based Cryptographic Accelerators for the Post-Quantum Era https://www.mdpi.com/2079-9292/15/2/475
16. [PDF] Hardware, Algorithms, and Applications of the Neuromorphic Vision ... https://arxiv.org/pdf/2504.08588?
17. Neuromorphic computing for robotic vision: algorithms to hardware ... https://pmc.ncbi.nlm.nih.gov/articles/PMC12350809/
18. Integration and Co-design of Memristive Devices and Algorithms for ... https://www.sciencedirect.com/science/article/pii/S2589004220310063
19. Integrated Neuromorphic Photonic Computing for AI Acceleration https://advanced.onlinelibrary.wiley.com/doi/10.1002/adma.202508029
20. Neuromorphic algorithms for brain implants: a review - Frontiers https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2025.1570104/full

21. Neuromorphic-based metaheuristics: A new generation of low ... https://arxiv.org/html/2505.16362v1
22. Neuromorphic Swarm on RRAM Compute-in-Memory Processor for ... https://ieeexplore.ieee.org/document/10247852/
23. On-demand photonic Ising machine with simplified Hamiltonian ... https://www.nature.com/articles/s42005-024-01658-x
24. Programmable photonic Ising machine enabled by a reconfigurable ... https://arxiv.org/html/2511.13284v1
25. Finding independent sets in large-scale graphs with a coherent Ising ... https://www.science.org/doi/10.1126/sciadv.ads7223
26. Single photon coherent Ising machines for constrained optimization ... https://iopscience.iop.org/article/10.1088/2058-9565/addde5/pdf
27. Efficient combinatorial optimization by quantum-inspired parallel ... https://www.nature.com/articles/s41467-023-41647-2
28. [PDF] Hardware solvers for combinatorial optimization problems - arXiv https://arxiv.org/pdf/2204.00276
29. Single photon coherent Ising machines for constrained optimization ... https://iopscience.iop.org/article/10.1088/2058-9565/addde5
30. ON-OFF neuromorphic ISING machines using Fowler-Nordheim ... https://pmc.ncbi.nlm.nih.gov/articles/PMC11958649/
31. [PDF] A Review of Ising Machines Implemented in Conventional and ... http://www.ece.ualberta.ca/~jhan8/publications/FINAL%20VERSION_Sept27.pdf
32. (PDF) Ising machines as hardware solvers of combinatorial ... https://www.researchgate.net/publication/359709550_Ising_machines_as_hardware_solvers_of_combinatorial_optimization_problems
33. [PDF] Training an Ising machine with equilibrium propagation - HAL https://hal.science/hal-04950942v1/file/s41467-024-46879-4.pdf
34. On polynomial selection for the general number field sieve https://www.researchgate.net/publication/220576812_On_polynomial_selection_for_the_general_number_field_sieve
35. Machine learning-aided differential-linear attacks with applications ... https://link.springer.com/article/10.1007/s44443-025-00241-w
36. [PDF] Deep Learning Assisted Key Recovery Attack for Round-Reduced ... https://bimsa.net/doc/publication/1545.pdf
37. [PDF] Improved (Related-key) Differential-based Neural Distinguishers for ... https://arxiv.org/pdf/2201.03767

38. Truncated Differential-Neural Key Recovery Attacks on Round ... https://www.mdpi.com/2079-9292/13/20/4053

39. (PDF) Truncated Differential-Neural Key Recovery Attacks on Round ... https://www.researchgate.net/publication/385113173_Truncated_Differential-Neural_Key_Recovery_Attacks_on_Round-Reduced_HIGHT

40. Enhanced related-key differential neural distinguishers for SIMON ... https://pmc.ncbi.nlm.nih.gov/articles/PMC11623040/

41. [PDF] 18.783 S2021 Lecture 11: Elliptic Curve Primality Proving (ECPP) https://ocw.mit.edu/courses/18-783-elliptic-curves-spring-2021/c95a8f3337131b100debe6d7a2c04b6e_MIT18_783S21_notes11.pdf

42. Deterministic elliptic curve primality proving for a special sequence ... https://arxiv.org/abs/1202.3695

43. (PDF) Prime Numbers:How We Need to Search for Large Numbers ... https://www.researchgate.net/publication/399669629_Prime_NumbersHow_We_Need_to_Search_for_Large_NumbersResonance_Lattice_Theory_of_Prime_Numbers_Transcendental_Geometry_and_Guided_Prime_Search

44. ON-OFF neuromorphic ISING machines using Fowler-Nordheim ... https://ideas.repec.org/a/nat/natcom/v16y2025i1d10.1038_s41467-025-58231-5.html

45. Benchmarking the optimization of optical machines with the planted ... https://www.nature.com/articles/s42005-024-01870-9

46. A versatile coherent Ising computing platform | Light - Nature https://www.nature.com/articles/s41377-025-02178-1

47. Programmable optoelectronic Ising machine for optimization of real ... https://www.researchgate.net/publication/399296193_Programmable_optoelectronic_Ising_machine_for_optimization_of_real-world_problems

48. ACE-SNN: Algorithm-Hardware Co-design of Energy-Efficient & Low ... https://pmc.ncbi.nlm.nih.gov/articles/PMC9025538/

49. A Quarter of a Century of Neuromorphic Architectures on FPGAs https://arxiv.org/html/2502.20415v2

50. Neuromorphic computing hardware and neural architectures for ... https://www.science.org/doi/10.1126/scirobotics.abl8419

51. Exploring Neuromorphic Computing Based on Spiking Neural ... https://dl.acm.org/doi/full/10.1145/3571155

52. (PDF) Analog optical computer for AI inference and combinatorial ... https://www.researchgate.net/publication/

395240933_Analog_optical_computer_for_AI_inference_and_combinatorial_optimization

53. Continuous-time digital twin with analog memristive neural ordinary ... https://www.science.org/doi/10.1126/sciadv.adr7571
54. Metrics for spin-based computing - arXiv https://arxiv.org/html/2510.17653v2
55. Top arXiv papers - SciRate https://scirate.com/?customerType=personal&date=2026-03-17&page=4&range=1
56. Mathematics, Volume 14, Issue 7 (April-1 2026) – 149 articles https://www.mdpi.com/2227-7390/14/7
57. Analog optical computer for AI inference and combinatorial ... https://www.microsoft.com/en-us/research/publication/analog-optical-computer-for-ai-inference-and-combinatorial-optimization/
58. Deep Learning Key Recovery for Simeck32/64 | PDF - Scribd https://www.scribd.com/document/980237406/1545
59. Efficient computation using spatial-photonic Ising machines with low ... https://www.nature.com/articles/s42005-025-01987-5
60. Higher-Order Neuromorphic Ising Machines - Autoencoders ... - arXiv https://arxiv.org/html/2506.19964v1
61. Higher-Order Neuromorphic Ising Machines -- Autoencoders and ... https://www.researchgate.net/publication/393022965_Higher-Order_Neuromorphic_Ising_Machines_--_Autoencoders_and_Fowler-Nordheim_Annealers_are_all_you_need_for_Scalability
62. ARTICLE IN PRESS - Nature https://www.nature.com/articles/s41467-026-71937-4_reference.pdf
63. Memristor-based hardware and algorithms for higher-order Hopfield ... https://www.researchgate.net/publication/381932257_Memristor-based_hardware_and_algorithms_for_higher-order_Hopfield_optimization_solver_outperforming_quadratic_Ising_machines
64. [PDF] Architecture Design for Analog Oscillatory Neural Networks https://theses.hal.science/tel-04561235v1/file/DELACOUR_2023_archivage.pdf
65. Appl. Sci., Volume 11, Issue 7 (April-1 2021) – 409 articles - MDPI https://www.mdpi.com/2076-3417/11/7
66. Algorithm-hardware co-design of neuromorphic networks with dual ... https://arxiv.org/abs/2512.07602
67. Algorithm Architecture Co-Design for Analog In-Memory Computing https://research.ibm.com/publications/algorithm-architecture-co-design-for-analog-in-memory-computing

68. Editorial: Algorithm-hardware co-optimization in neuromorphic ... https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2025.1746610/full
69. Algorithm-hardware co-design of binary neural network for efficient ... https://www.researchgate.net/publication/400419614_Algorithm-hardware_co-design_of_binary_neural_network_for_efficient_super_resolution_on_FPGA
70. Device-Algorithm Co-Design of Ferroelectric Compute-in-Memory In ... https://www.researchgate.net/publication/395629599_Device-Algorithm_Co-Design_of_Ferroelectric_Compute-in-Memory_In-Situ_Annealer_for_Combinatorial_Optimization_Problems
71. Future-Proofing Cloud Security Against Quantum Attacks - arXiv https://arxiv.org/html/2509.15653v1
72. High-Performance Ideal Lattice-Based Cryptography on 8-Bit ... https://www.researchgate.net/publication/283842736_High-Performance_Ideal_Lattice-Based_Cryptography_on_8-Bit_ATxmega_Microcontrollers
73. [PDF] Securing Cloud Computing Against Quantum Threats - arXiv https://arxiv.org/pdf/2509.15653
74. Applications in Electronics Pervading Industry, Environment and ... https://link.springer.com/content/pdf/10.1007/978-3-031-30333-3.pdf
75. Polynomial inversion hardware accelerator for post quantum algorithm https://www.researchgate.net/publication/395368405_Polynomial_inversion_hardware_accelerator_for_post_quantum_algorithm
76. Machine learning-aided differential-linear attacks with applications ... https://www.researchgate.net/publication/395796643_Machine_learning-aided_differential-linear_attacks_with_applications_to_Des_and_Speck3264
77. (PDF) Elliptic Curves And Primality Proving - ResearchGate https://www.researchgate.net/publication/2422617_Elliptic_Curves_And_Primality_Proving
78. [PDF] arXiv:math/0502097v1 [math.NT] 4 Feb 2005 https://arxiv.org/pdf/math/0502097
79. Informed Search Algorithms in Artificial Intelligence - GeeksforGeeks https://www.geeksforgeeks.org/artificial-intelligence/informed-search-algorithms-in-artificial-intelligence/
80. AI Search Strategies Overview | PDF | Applied Mathematics - Scribd https://www.scribd.com/document/884204359/Search-Strategies-Algorithms-and-Techniques-in-AI
81. [PDF] arXiv:2211.13316v4 [cs.AI] 16 Feb 2025 https://arxiv.org/pdf/2211.13316
82. (PDF) An On-demand Photonic Ising Machine with Simplified ... https://www.researchgate.net/publication/370517247_An_On-

demand_Photonic_Ising_Machine_with_Simplified_Hamiltonian_Calculation_by_Phase-encoding_and_Intensity_Detection

83. Photonic Spatial-Euler Ising Machine for Solving 20000-node Max ... https://www.researchgate.net/publication/367088422_Photonic_Spatial-Euler_Ising_Machine_for_Solving_20000-node_Max-cut_Problem
84. Higher-Order Neuromorphic Ising Machines - LinkedIn https://www.linkedin.com/posts/shantanu-chakrabartty-6563754_higher-order-neuromorphic-ising-machines-activity-7352502308124971008-YGau
85. A Quarter of a Century of Neuromorphic Architectures on FPGAs https://arxiv.org/html/2502.20415v3
86. The neurobench framework for benchmarking neuromorphic ... https://www.nature.com/articles/s41467-025-56739-4
87. [PDF] FPGA-Based Neuromorphic Architecture for Spiking Neural Network ... https://hal.science/tel-05543829v1/file/These.pdf
88. Memristor in a Reservoir System—Experimental Evidence for High ... https://pubs.acs.org/doi/10.1021/acsami.9b01841
89. 2022 roadmap on neuromorphic devices and applications research ... https://iopscience.iop.org/article/10.1088/2634-4386/ac7a5a
90. On-device Online Learning and Semantic Management of TinyML ... https://dl.acm.org/doi/10.1145/3665278
91. A Deeper Look at Machine Learning-Based Cryptanalysis https://dl.acm.org/doi/abs/10.1007/978-3-030-77870-5_28
92. 人工智能2026_4_22 - arXiv每日学术速递 https://www.arxivdaily.com/thread/79084
93. Electronics, Volume 15, Issue 5 (March-1 2026) – 232 articles https://www.mdpi.com/2079-9292/15/5
94. Intelligent neuromorphic computing based on nanophotonics and ... https://www.researchgate.net/publication/378081827_Intelligent_neuromorphic_computing_based_on_nanophotonics_and_metamaterials
95. [PDF] Lattice Reduction Algorithms - Semantic Scholar https://www.semanticscholar.org/paper/Lattice-Reduction-Algorithms-Stehl%C3%A9/95c553b3a7d36019ab27dc4c539cbe86203b7614
96. Issue 9 - Volume 100 - Physica Scripta - IOPscience https://iopscience.iop.org/issue/1402-4896/100/9
97. Mathematics, Volume 13, Issue 17 (September-1 2025) – 208 articles https://www.mdpi.com/2227-7390/13/17

98. The 35th Chinese Control and Decision Conference - IEEE Xplore https://ieeexplore.ieee.org/iel7/10326447/10326467/10327206.pdf
99. Mathematics Jul 2024 - arXiv http://arxiv.org/list/math/2024-07?skip=3570&show=1000
100. The future of electronics based on memristive systems - ResearchGate https://www.researchgate.net/publication/322316132_The_future_of_electronics_based_on_memristive_systems
101. [XLS] China https://www.nsfc.gov.cn/Portals/0/fj/fj20230220_01.xlsx
102. Lattice-Based Cryptographic Accelerators for the Post-Quantum Era https://www.researchgate.net/publication/400038783_Lattice-Based_Cryptographic_Accelerators_for_the_Post-Quantum_Era_Architectures_Optimizations_and_Implementation_Challenges
103. [PDF] arXiv:2502.18570v1 [quant-ph] 25 Feb 2025 https://arxiv.org/pdf/2502.18570
104. Fully Programmable Spatial Photonic Ising Machine by Focal Plane ... https://www.researchgate.net/publication/389003895_Fully_Programmable_Spatial_Photonic_Ising_Machine_by_Focal_Plane_Division
105. Jason K. Eshraghian University of California, Santa Cruz https://www.researchgate.net/profile/Jason-Eshraghian
106. Machine learning algorithms for FPGA Implementation in biomedical ... https://www.sciencedirect.com/science/article/pii/S2405844024026835
107. [PDF] A Unified Framework for Combinatorial Optimization Based ... - arXiv https://arxiv.org/pdf/2406.13125
108. Towards Next-Generation FPGA-Accelerated Vision-Based ... - MDPI https://www.mdpi.com/2624-6120/6/4/53
109. Application of Meta-Heuristic Algorithms for Training Neural ... - PMC https://pmc.ncbi.nlm.nih.gov/articles/PMC9628382/
110. Simulation of Quantum Computers: Review and Acceleration ... https://dl.acm.org/doi/full/10.1145/3762672
111. LIMO: Low-power in-memory-annealer and matrix-multiplication ... https://www.nature.com/articles/s44335-026-00054-8
112. A comprehensive benchmark of an Ising machine on the Max-Cut ... https://arxiv.org/abs/2507.22117
113. [PDF] Optoelectronic coherent ising machine for combinatorial ... - HAL https://hal.science/hal-04068504/file/Optics_Letter_CIM.pdf
114. Experimental investigation of performance differences between ... https://pmc.ncbi.nlm.nih.gov/articles/PMC6534389/

115. Solving MAXCUT Optimization Problems with a Coherent Ising ... https://www.semanticscholar.org/paper/c93743158e458b415daa39c33fd5a00a3b3890e8
116. ECM on Graphics Cards | Request PDF - ResearchGate https://www.researchgate.net/publication/220334340_ECM_on_Graphics_Cards
117. [PDF] Practical Lightweight Security - HAL https://hal.science/tel-05200759v1/file/Dissertation_Anagnostopoulos_2022_two-sided.pdf
118. Issue 02 - Volume 18 - Journal of Instrumentation - IOPscience https://iopscience.iop.org/issue/1748-0221/18/02
119. Sensors, Volume 21, Issue 6 (March-2 2021) – 328 articles - MDPI https://www.mdpi.com/1424-8220/21/6?listby=date&view=default&section_id=11
120. Latest articles - MDPI https://www.mdpi.com/latest_articles
121. Technologies, Volume 13, Issue 10 (October 2025) – 53 articles https://www.mdpi.com/2227-7080/13/10
122. AI and ML empowering 5G and shaping the 6G future http://www.sciencedirect.com/science/article/pii/S2405959525001985;
123. (PDF) Review of deep learning: concepts, CNN architectures ... https://www.researchgate.net/publication/350527503_Review_of_deep_learning_concepts_CNN_architectures_challenges_applications_future_directions
124. [PDF] Testability and Dependability of AI Hardware: Survey, Trends ... - HAL https://hal.science/hal-03961502/file/Testability_and_Dependability_of_AI_Hardware__Survey__Trends__Challenges__and_Perspectives.pdf
125. Papers For Practitioners - ACM https://www.acm.org/publications/papers-for-practitioners