Optoelectronic coherent ising machine for combinatorial optimization problems

Nickson Mwamsojo, Frederic Lehmann, Kamel Merghem, Badr-Eddine Benkelfat, Yann Frignac

Optoelectronic coherent Ising machine for combinatorial optimization problems

NICKSON MWAMSOJO[1,*], FREDERIC LEHMANN[1], KAMEL MERGHEM[1], BADR-EDDINE BENKELFAT[1], AND YANN FRIGNAC[2]

[1] TIPIC - SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, 91120 Palaiseau, France
[2] Huawei Technologies France, 18 Quai du Point du Jour, 92100 Boulogne-Billancourt, France
[*] Corresponding author: nickson_mwamsojo@telecom-sudparis.eu

Compiled March 29, 2023

Hopfield networks are iterative procedures able to solve combinatorial optimization problems. New studies regarding algorithm-architecture adequacy are fostered by the re-emergence of hardware implementations of such methods in the form of Ising machines. In this work, we propose an optoelectronic architecture suitable for fast processing and low energy consumption. We show that our approach allows effective optimization relevant to statistical image denoising.

http://dx.doi.org/10.1364/ao.XX.XXXXXX

Bio-inspired physical systems made from a large number of interconnected simple elements can yield interesting computational properties. In his seminal work, Hopfield introduced networks of neuron-like elements and their usefulness in combinatorial optimization problems expressed as the minimization of a quadratic energy function, both in continuous [1] and discrete time [2]. While theoretical guidelines to achieve feasible solutions were investigated in [3], a number of practical improvements have also been published to escape local minima including tailoring the energy function and its hyperparameters [4], incorporating simulated annealing heuristics [5], random noise [6] or transient chaotic behavior [7]. A generalization to highly nonlinear energy functions has also appeared in [8].

The introduction of efficient hardware architectures of Hopfield networks and their generalizations during the last decade has sparked renewed interest in this field. These implementations, known as Ising machines, are based on various physical principles such as quantum [9], nanomagnetic [10], memresistive [11] and photonic [12] technologies. Note that specific methods for optical Ising machines to escape local minima, have also appeared recently in [13]-[14]. Unfortunately, most papers on the subject lack a comparison with respect to near-global optimization techniques as well as an evaluation of energy efficiency.

In this letter, we implement a mixed hardware/digital architecture of a coherent Ising machine (CIM) based on off-the-shelf telecommunication components. The main novelties are as follows: from an application perspective, we demonstrate the potential of a CIM for statistical image denoising - from a system-level perspective, we compare the proposed mixed hardware/digital system to a standard digital implementation of Hopfield networks and Simulated Annealing (SAN) [15] in terms of probability of reaching the ground state, computational complexity, and energy consumption.

We implement the coherent Ising model in a digital simulation running on a CPU and the equivalent optoelectronic hardware. We describe the model formally by letting Ω be a square $n \times n$ lattice. Any pixel $s = (l, c) \in \Omega$ can be assimilated to a position in the lattice with line (resp. column) coordinate l (resp. c), where $1 \leq l, c \leq n$. $\forall s \in \Omega$, we let the spin σ_s be a random variable in $\{-1, +1\}$. The random vector $\boldsymbol{\sigma}$ is obtained by raster scanning the spins columnwise. A particular Markov random field (MRF) defines the probability mass function (pmf) of $\boldsymbol{\sigma}$ as

$$P(\boldsymbol{\sigma} = (\sigma_s)_{s\in\Omega}) \propto e^{-E(\boldsymbol{\sigma})}, \tag{1}$$

where the system's energy function has the form of an Ising Hamiltonian

$$E(\boldsymbol{\sigma}) = -\frac{1}{2}\sum_{(s,t)\in\mathcal{N}} J_{s,t}\sigma_s\sigma_t - \sum_{s\in\Omega} b_s\sigma_s, \tag{2}$$

where $\mathcal{N}$ denotes all couples of neighboring pixels with end-around boundary conditions. Note that finding the most probable spin configuration (the so-called ground state) is equivalent to the combinatorial optimization problem consisting in minimizing Equation 2.

The numerical simulation follows [12] by numerically implementing a generalized Hopfield network to minimize the discrete time equation of the form:

$$\begin{aligned} x_s(k) &= f\left(\alpha x_s(k-1) + \beta\left(\sum_{t:(s,t)\in\mathcal{N}} J_{s,t}x_t(k-1) + b_s\right)\right) \\ \hat{\sigma}_s(k) &= sign(x_s(k)), \quad \forall s \in \Omega \end{aligned} \tag{3}$$

where k is the discrete-time, $f(.)$ is a nonlinear activation function, α and β are scaling coefficients controlling the self-coupling and feedback strength affecting the neuron output $x_s(k)$, while $\hat{\sigma}_s(k)$ is the spin estimate of pixel s. We call *digital CIM* a Python implementation of the algorithm described by the equation above executed on an Intel Xeon E5-1603 processor. The nonlinear function is the $sin^2(.)$ imitating the transfer function of the MZM followed by a photodiode used in the hardware implementation.

Building on the digital implementation we implement mixed analog/digital processing in the form of a CIM that can lead to savings in terms of achievable processing speeds and energy consumption. The potential for high-speed solutions to combinatorial optimization problems is very attractive for applications requiring real-time processing or adaptation to dynamically changing environments. We solve the algorithm-architecture adequacy problem by selecting only commercially available telecommunication components in the optoelectronic oscillator setup shown in Figure 1. We call this setup the *hardware CIM*.

The hardware CIM uses a Distributed Feedback (DFB) laser diode emitting light at $1.55\mu m$ to a Mach-Zehnder Modulator (MZM) that modulates the optical phase of the light. The phase is modulated proportionally to the feedback signal allowing for the electro-optical conversion of the system's state. The MZM also provides the nonlinear activation function in the optical domain thanks to its $sin^2(.)$ transfer function. Note that alternative nonlinearities could improve the photonic optimizer [16], but at the expense of more complex hardware. The modulated light from the MZM is photodetected by a 20 GHz photodiode and the electrical signal is sent to the Zmod Scope $1410-125$ Analog-to-Digital Converter with $125MSa/s$ and 14-bit resolution. The scope supplies the digital part with the node-states vector for the spin interactions computation according to Equation 3 and the subsequent digital processing of the spins. The digital part is an Eclypse-Z7 featuring the Field Programmable Gate Array (FPGA) board along with two ZMod connectors allowing high-speed data conversions and processing. We implement via VHSIC Hardware Description Language (VHDL) the logic for high-speed data acquisition, matrix multiplication, and transfer of signals. After digital processing, the resultant node-states vector is sent to a Zmod AWG 1411: 2-channel 14-bit Arbitrary Waveform Generator (AWG) with a sample rate of $100MSa/s$.This allows for the conversion of the FPGA digital data to a continuous signal that modulates the phase of the MZM completing the loop. We also incorporate a 10 GHz Analog RF Amplifier Driver with an output voltage of $9V_{pp}$ allowing proper signal scaling before modulating the MZM and a Digital Serial Analyser (DSA) to visualize in real-time the evolution of the signals and spin formation in the loop.

The presented setup can implement Equation 3 only for a single spin, the feedback delay is decomposed into n^2 intervals along which the spins are multiplexed using Time Division Multiplexing (TDM). For each iteration, the FPGA waits for n^2 readings from the ADC before computing the resultant spins after spin interaction in a way that is more reliable, scalable and inexpensive than today's fully-photonic alternatives [17].

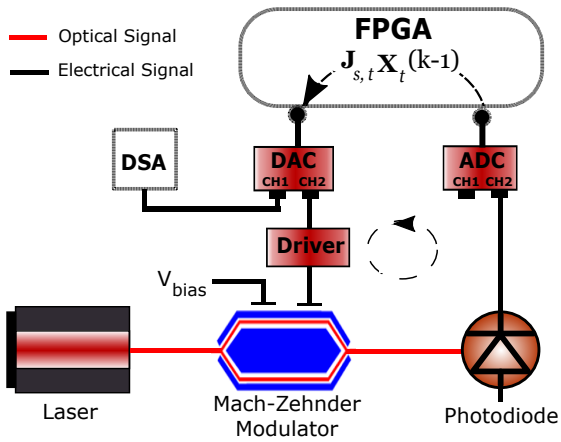

Fig. 1. Experimental setup of the proposed opto-electronic CIM. V_{bias} is the bias voltage control of the MZM.

As a benchmark, we implemented SAN similar to [18], arguably the most used heuristic method for gradual *cooling* of a 'high-temperature' problem to attain a frozen state that is, ideally, arbitrarily close to the solution of the problem. SAN is one of the most popular algorithms for this feat and has been implemented both in computer simulations as well as in dedicated hardware which provides parallelization of digital hardware accelerators and analog computing [19].

To compare the optimization implementations under study we propose several metrics. Firstly, we evaluate the success probability of reaching (or approximating) the ground state for each method and compare them. We extend the analysis to incorporate the estimation of computational complexity, that in the digital domain, is measured via the number of Floating Point Operations (FPOs), while an *experiment-impact-tracker* [20] running in parallel is used to report the consumed energy and the corresponding carbon-dioxide and equivalent greenhouse gases emissions (CO_{2eq}). For the hardware CIM, we estimate the average power consumption of constituent elements to obtain a rough estimate of the consumption. These metrics enable further analysis of the studied approaches by shedding light on the complexity vs. energy vs. performance trade-off, a subject that is mostly overlooked in the literature despite the recommendations for sustainable practices [21].

The proposed performance metrics gauge the ability and efficiency of the digital CIM and hardware CIM. We also propose the near-global optimization using SAN based on the Gibbs sampler given in [22] as the benchmark. Unless otherwise stated, $\mathcal{N}$ is chosen as the couple of 4-point nearest neighbors in the lattice under end-around boundary conditions. For the digital and optoelectronic CIM, the initial spin configuration is chosen uniformly and independently at random. The hyperparameters are set to $(\alpha = 0.25, \beta = 0.29)$ (value retained in [12] Fig. 4(a) for 2D square lattices), while the number of iterations is set to $N_{it} = 100$. Also, we implement SAN with an initial temperature of 2, a geometric annealing parameter equal to 0.99, and 200 iterations [22]. Moreover, the adopted legend for all images is as follows: bright yellow for spin-up $(+1)$ and dark purple for spin-down (-1). We study two problems that can be formulated as the minimization of a Hamiltonian having the form given by Equation 2 such as the Antiferromagnetic Ising model and MAP image denoising using the optoelectronic Coherent Ising Machine (CIM) both in the simulations and hardware.

We begin with a low-dimensional example where $n = 10$ with antiferromagnetic interactions, that is $J_{s,t} = -1\ \forall (s,t) \in \mathcal{N}$ and $b_s = 0\ \forall s \in \Omega$. It is well-known that the ground state corresponds to a checkerboard pattern (alternating spin-up and spin-down configuration with minimum energy equal to $-n^2$). In our experiments, Figure 2 depicts the initial and final estimated spin configuration for a single run of the proposed optoelectronic CIM, that successfully converges to the ground state. Part of the TDM sequence of spins is copied on channel CH1 of DAC and sent to a Digital Serial Analyzer Sampling Oscilloscope (DSA). A screenshot in Figure 3 shows the resultant alternating up and down spins as a time series. Over a single successful run, Figure 4 shows that the proposed optoelectronic CIM and the digital CIM have similar dynamics, with convergence reached typically after a few tens of iterations. Table 1 summarizes our perfor-

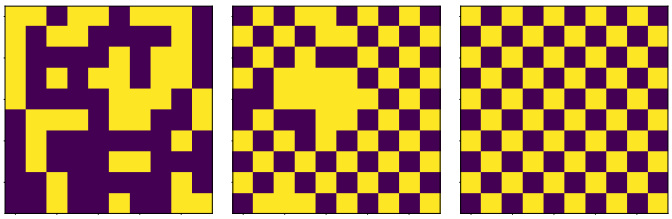

(a) Iteration = 1 **(b)** Iteration = 25 **(c)** Iteration > 50

Fig. 2. The initial spins (Iteration = 0) are randomly and independently chosen for the digital CIM whereas, for the optoelectronic CIM, the system's noise initializes the spins. A checkerboard pattern appears (Iteration = 25) and stabilizes as the system converges (Iteration = 50 to 200).

mance metrics by repeating all the aforementioned experiments independently 1000 times.

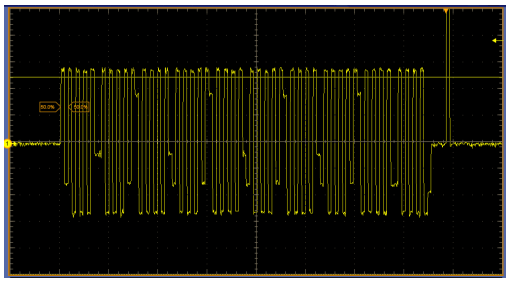

Fig. 3. Digital Serial Analyzer screen showing serial alternating spins after convergence for the hardware CIM - $n = 10$.

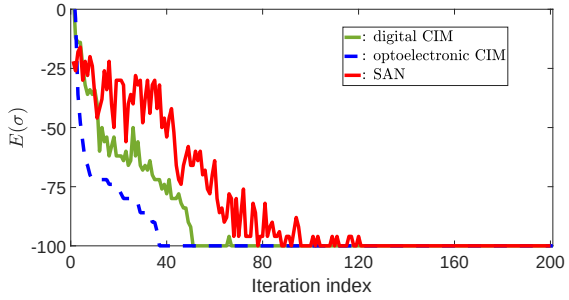

Fig. 4. Energy evolution for a single run for the square-lattice of spins with antiferromagnetic interactions - $n = 10$.

We observe that SAN converges later than both CIMs but reaches the ground state with approximately 98.9% success probability compared to the approximately 91.4% digital CIM. For this task, success means attaining exactly the theoretical ground energy state of $E(\sigma) = -100$. From this standpoint, SAN is more performant. With further analysis, we noticed, however, that this win comes at a cost of 7.47 times the number of FPOs and runtime required by the digital CIM. It, therefore, takes more computational effort to attain a solution with SAN than it does with digital CIM. Moreover, energy estimates with the *experiment-impact-tracker* show that SAN requires 23 times the energy of the digital CIM and the same factor for the increase in

Metric	Digital CIM	Optoeletronic CIM	SAN
Success probability (%)	91.4	90	98.9
Time (ms)	10	2.1	79
FPOs ($\times 10^6$)	0.76	0.32	5.68
Energy ($\times 10^{-6}$ kWh)	0.6	3×10^{-4}	14
CO_{2eq} ($\times 10^{-6} kg$)	0.044	2.2×10^{-5}	1

Table 1. Average performance metrics for the antiferromagnetic model over 1000 runs - $n = 10$.

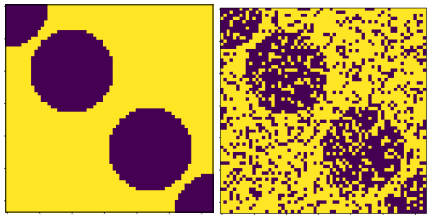

(a) **(b)**

Fig. 5. The clean image (a) and the resultant noisy image (b) after the salt and pepper noise addition - $n = 64$, $p = 0.2$.

CO_{2eq} emissions. The hardware CIM has a success probability of 90% - close to that of the digital CIM. Since the hardware CIM benefits from the speed of optics and the FPGA's programmable logic the energy analysis for the hardware CIM becomes interesting. We observe that the checkerboard solution obtained in digital CIM requires 2000 times the cost of the optoelectronic CIM counterpart. What's more, the hardware CIM consumes $1/46667^{th}$ of the energy required by SAN. A significant gain in efficiency altogether.

We now consider a hidden black and white image $(\sigma_s)_{s\in\Omega}$ to be restored from an observed noisy image $(y_s)_{s\in\Omega}$ according to a salt and pepper noise model, i.e. $y_s = -\sigma_s$ with probability p and $y_s = \sigma_s$ with probability $1-p$, independently for each pixel $s \in \Omega$. In our setting, $n = 64$ and the prior spin pmf is chosen as the MRF in Equation 1 with the coupling parameter between neighboring pixels set to 1 (ferromagnetic interactions). It is easily shown that maximum a posteriori (MAP) image denoising corresponds to selecting $J_{s,t} = 1\ \forall (s,t) \in \mathcal{N}$ and $b_s = -\frac{1}{2}\log(p/(1-p))y_s\ \forall s \in \Omega$ (as derived in supplementary material). In our experiments, we use $p = 0.2$ and we show the clean image and the resultant noisy image after being impacted by the noise in Figure 5. Note that further results (although not reported) showed the robustness of the proposed denoising against different image structure and values for n and p.

Figure 6 depicts the initial dirty image and the final image after running SAN and digital/optoelectronic CIM. Over a single successful run, Figure 7 shows that the proposed optoelectronic CIM and the digital CIM have similar dynamics, with convergence reached typically after 15 iterations. Table 2 summarizes our performance metrics - adding the pixel-wise classification error rate (PCER %) - by repeating all aforementioned experiments independently 1000 times.

For this application, we observe that SAN attained the ground state successfully for all the runs whereas digital and hardware CIM attained the ground state in 89% and 86.75% of the runs respectively. Success means falling within three standard deviations of the average ground state energy (which is approximately

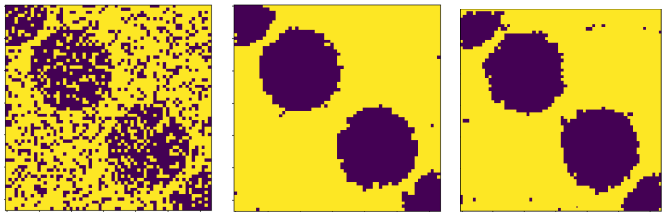

(a) (b) (c)

Fig. 6. A sample initial dirty image (iteration 1) shown in (a). After convergence (iteration 100), we obtain (b) and (c) for the digital and optoelectronic CIMs respectively - $n = 64$, $p = 0.2$.

$E(\sigma) = -5550$), thus corresponding to a 98% confidence interval. The clean images generated by SAN contain a PCER of 1.7% whereas the digital and hardware CIMs converged to a PCER of 2.3% and 3% respectively. For this task, SAN excels over the digital and hardware CIMs in these convergence metrics also. Nevertheless, further analysis reveals that SAN's performance comes at approximately 9.5 times the execution time, 10 times the number of FPOs, and 12.22 times the energy (same factor for the CO_{2eq} emissions) consumed by the digital CIM. The CO_{2eq} emissions are reported taking into account the nature of electric grids in the *Palaiseau* city in *France*. The 11% gain in convergence success probability of SAN costs us at least 10 times the computation cost of digital CIM by all measures. Following the observation, with the Antiferromagnetic model we analyze the energy costs for the hardware CIM on this task as well. The energy consumption results for hardware CIM are reported in Table 2 showing a gain factor of 6206 and 75862 with respect to the digital CIM and SAN respectively for the energy and environmental impact.

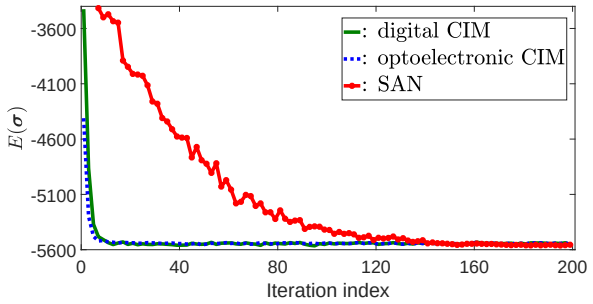

Fig. 7. MAP image denoising energy evolution for a single run - $n = 64$, $p = 0.2$.

Metric	Digital CIM	Optoeletronic CIM	SAN
Success probability (%)	89	86.75	100
PCER (%)	2.3	3	1.7
Time (ms)	4110	86	39130
FPOs ($\times 10^9$)	0.7	0.0131	7.1
Energy ($\times 10^{-6}$ kWh)	90	1.45×10^{-2}	1100
CO_{2eq} ($\times 10^{-6} kg$)	5	8×10^{-4}	62

Table 2. Average performance metrics for MAP image denoising over 1000 runs - $n = 64$, $p = 0.2$.

In summary, we have proposed and demonstrated the potential of CIM for image restoration applications. We have also illustrated that SAN outperforms our CIMs in terms of convergence success probability and PCER. We have argued, however, that SAN requires significantly more resources compared to digital and hardware CIM implementations. With the severalfold increase in energy costs and computation efforts, a compromise on accuracy becomes reasonable. In most practical applications, a choice is often made to settle for less demanding solutions that yield acceptable performances. In this light, CIMs appear as the more reasonable and informed choice. And further energy efficiency is attained with an optoelectronic hardware CIM compared to the digital one, also by a significant margin. This underlines the argument for the efficiency of neuromorphic photonic implementations of brain-like algorithms. As we explore more complex combinatorial optimization problems and their interesting applications in various domains we will be mindful of the CIM's promise for faster and more efficient solutions, this direction of exploration has the potential for interesting avenues.

REFERENCES

1. J. Hopfield, National Acad. Sci. United States Am. **81**, 3088 (1984).
2. J. Hopfield and D. Tank, Biol. cybernetics **52**, 141 (1985).
3. S. Aiyer, M. Niranjan, and F. Fallside, IEEE Trans. on Neural Networks **1**, 204 (1990).
4. J.-Y. Potvin, INFORMS J. Comput. **5**, 328 (1993).
5. D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, Cogn. Sci. **9**, 147 (1985).
6. Y. Akiyama, A. Yamashita, M. Kajiura, and H. Aiso, "Combinatorial optimization with gaussian machines," in *International 1989 Joint Conference on Neural Networks*, , vol. 1 (Washington, DC, USA, 1989), pp. 533–540.
7. L. Chen and K. Aihara, Neural Networks **8**, 915 (1995).
8. D. Krotov, arXiv pp. 1–13 (2021).
9. E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, Science **292**, 472 (2001).
10. B. Sutton, K. Y. Camsari, B. Behin-Aein, and S. Datta, Sci. Reports **7**, 1 (2017).
11. Z. Fahimi, M. R. Mahmoodi, H. Nili, V. Polishchuk, and D. B. Strukov, Sci. Reports **11**, 16383 (2021).
12. F. Böhm, G. Verschaffelt, and G. Van der Sande, Nat. Commun. **10**, 3538 (2019).
13. T. Leleu, Y. Yamamoto, P. L. McMahon, and K. Aihara, Phys. Rev. Lett. **122**, 040607 (2019).
14. B. Lu, C.-R. Fan, L. Liu, K. Wen, and C. Wang, Opt. Express **31**, 3676 (2023).
15. J. C. Spall, *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control* (John Wiley & Sons, Ltd, 2003), chap. Stochastic Search and Optimization: Motivation and Supporting Results, pp. 1–33.
16. F. Böhm, T. V. Vaerenbergh, G. Verschaffelt, and G. V. der Sande, Commun. Phys. **4**, 1 (2021).
17. M. Prabhu, C. Roques-Carmes, Y. Shen, N. Harris, L. Jing, J. Carolan, R. Hamerly, T. Baehr-Jones, M. Hochberg, V. Čeperić *et al.*, Optica **7**, 551 (2020).
18. S. Geman and D. Geman, IEEE Trans. on Pattern Anal. Mach. Intell. **6**, 721 (1984).
19. N. Mohseni, P. L. McMahon, and T. Byrnes, Nat. Rev. Phys. **4**, 363–379 (2022).
20. P. Henderson, J. Hu, J. Romoff, E. Brunskill, D. Jurafsky, and J. Pineau, J. Mach. Learn. Res. **21**, 10039–10081 (2020).
21. R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, Commun. ACM (CACM) **63**, 54 (2020).
22. J. Besag, J. royal statistical society series b-methodological **48**, 259 (1986).

FULL REFERENCES

1. J. Hopfield, "Neurons with graded response have collective computational properties like those of two-state neurons," National Acad. Sci. United States Am. **81**, 3088–92 (1984).
2. J. Hopfield and D. Tank, "Neural computation of decisions in optimization problems," Biol. cybernetics **52**, 141–52 (1985).
3. S. Aiyer, M. Niranjan, and F. Fallside, "A theoretical investigation into the performance of the hopfield model," IEEE Trans. on Neural Networks **1**, 204–215 (1990).
4. J.-Y. Potvin, "State-of-the-art survey - the traveling salesman problem: A neural network perspective," INFORMS J. Comput. **5**, 328–348 (1993).
5. D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, "A learning algorithm for boltzmann machines," Cogn. Sci. **9**, 147–169 (1985).
6. Y. Akiyama, A. Yamashita, M. Kajiura, and H. Aiso, "Combinatorial optimization with gaussian machines," in *International 1989 Joint Conference on Neural Networks,* vol. 1 (Washington, DC, USA, 1989), pp. 533–540.
7. L. Chen and K. Aihara, "Chaotic simulated annealing by a neural network model with transient chaos," Neural Networks **8**, 915–930 (1995).
8. D. Krotov, "Hierarchical associative memory," arXiv pp. 1–13 (2021).
9. E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, "A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem," Science **292**, 472–475 (2001).
10. B. Sutton, K. Y. Camsari, B. Behin-Aein, and S. Datta, "Intrinsic optimization using stochastic nanomagnets," Sci. Reports **7**, 1–15 (2017).
11. Z. Fahimi, M. R. Mahmoodi, H. Nili, V. Polishchuk, and D. B. Strukov, "Combinatorial optimization by weight annealing in memristive hopfield networks," Sci. Reports **11**, 16383 (2021).
12. F. Böhm, G. Verschaffelt, and G. Van der Sande, "A poor man's coherent ising machine based on optoelectronic feedback systems for solving optimization problems," Nat. Commun. **10**, 3538 (2019).
13. T. Leleu, Y. Yamamoto, P. L. McMahon, and K. Aihara, "Destabilization of local minima in analog spin systems by correction of amplitude heterogeneity," Phys. Rev. Lett. **122**, 040607 (2019).
14. B. Lu, C.-R. Fan, L. Liu, K. Wen, and C. Wang, "Speed-up coherent ising machine with a spiking neural network," Opt. Express **31**, 3676–3684 (2023).
15. J. C. Spall, *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control* (John Wiley & Sons, Ltd, 2003), chap. Stochastic Search and Optimization: Motivation and Supporting Results, pp. 1–33.
16. F. Böhm, T. V. Vaerenbergh, G. Verschaffelt, and G. V. der Sande, "Order-of-magnitude differences in computational performance of analog ising machines induced by the choice of nonlinearity," Commun. Phys. **4**, 1–11 (2021).
17. M. Prabhu, C. Roques-Carmes, Y. Shen, N. Harris, L. Jing, J. Carolan, R. Hamerly, T. Baehr-Jones, M. Hochberg, V. Čeperić *et al.*, "Accelerating recurrent ising machines in photonic integrated circuits," Optica **7**, 551–558 (2020).
18. S. Geman and D. Geman, "Stochastic relaxation, gibbs distributions, and the bayesian restoration of images," IEEE Trans. on Pattern Anal. Mach. Intell. **6**, 721–741 (1984).
19. N. Mohseni, P. L. McMahon, and T. Byrnes, "Ising machines as hardware solvers of combinatorial optimization problems," Nat. Rev. Phys. **4**, 363–379 (2022).
20. P. Henderson, J. Hu, J. Romoff, E. Brunskill, D. Jurafsky, and J. Pineau, "Towards the systematic reporting of the energy and carbon footprints of machine learning," J. Mach. Learn. Res. **21**, 10039–10081 (2020).
21. R. Schwartz, J. Dodge, N. A. Smith, and O. Etzioni, "Green AI," Commun. ACM (CACM) **63**, 54–63 (2020).
22. J. Besag, "On the statistical analysis of dirty pictures," J. royal statistical society series b-methodological **48**, 259–279 (1986).