





Genetic-tunneling driven energy optimizer for spin systems

Qichen Xu ^{1,2}✉, Zhuanglin Shen³, Manuel Pereiro ⁴, Erik Sjöqvist⁴, Pawel Herman ^{2,5}, Olle Eriksson⁴ & Anna Delin ^{1,2}✉

Finding the ground state of complex many-body systems, such as magnetic materials containing topological textures, like skyrmions, is a fundamental and long-standing problem. We present here a genetic-tunneling-driven variance-controlled optimization method, that efficiently identifies the ground state of two-dimensional skyrmionic systems. The approach combines a local energy-minimizer backend and a metaheuristic global search frontend. The method is shown to perform significantly better than simulated annealing. Specifically, we demonstrate that for the Pd/Fe/Ir(111) system, our method correctly and efficiently identifies the experimentally observed spin spiral geometry, skyrmion lattice and ferromagnetic ground states as a function of the external magnetic field. To our knowledge, no other optimization method has until now succeeded in doing this. We envision that our findings will pave the way for evolutionary computing in mapping out phase diagrams for spin systems in general.

¹Department of Applied Physics, School of Engineering Sciences, KTH Royal Institute of Technology, AlbaNova University Center, SE-10691 Stockholm, Sweden. ²SeRC (Swedish e-Science Research Center), KTH Royal Institute of Technology, SE-10044 Stockholm, Sweden. ³CAS Key Laboratory of Quantitative Engineering Biology, Shenzhen Institute of Synthetic Biology, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, Shenzhen 518055, China. ⁴Department of Physics and Astronomy, Uppsala University, Box 516, SE-75120 Uppsala, Sweden. ⁵Division of Computational Science and Technology, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, AlbaNova University Center, SE-10691 Stockholm, Sweden. ✉email: qichenx@kth.se; annadel@kth.se

Optimization algorithms are central in many areas of physics, for instance whenever one is dealing with complex systems such as the many-body problem. A common challenge is to find the global minimum in the potential energy surface (PES) describing the system. The PES is often extremely complicated with numerous local minima, making it very hard to identify the global minimum. Interestingly, these types of systems can often be described using the language of spin models. Therefore, spin models have widely used in areas not directly connected to magnetism, e.g., percolation theory, protein folding, and stock market trading models¹.

In this work, we specifically focus on the important class of two-dimensional (2D) spin systems with complex magnetic interactions—interactions that give rise to frustration or convoluted magnetic textures, such as for instance, skyrmions or other magnetic topological structures^{2–4}. For such systems, a number of optimization algorithms have been developed, e.g., gradient-descent-based methods, Monte Carlo approaches, and methods based on spin dynamics^{5–7}. A recurring problem, however, is the tendency to get trapped into one of the many local energy minima in the PES (the freezing problem) rather than converging toward the global minimum within a reasonable amount of time.

To attempt to improve on this situation, a viable route is to explore meta-heuristic optimization methods aimed at efficiently exploring the search space in order to find near-optimal solutions. In particular, Markov chain Monte Carlo (MCMC) based heat-bath methods—a group of non-gradient sampling algorithms—have proven themselves to be both effective and robust regarding searching for low-energy states at finite temperature^{8,9}. Unfortunately, in their current implementations^{10,11}, there is still a great deal of prior knowledge needed—e.g., appropriate initial guesses and manual convergence analysis—in order to achieve acceptable results.

In order to remedy these shortcomings, several hybrid approaches based on the idea of combining metaheuristic algorithms and typical optimization approaches, e.g., hybrid Monte Carlo¹², neural annealing optimization¹³ or neural evolutionary methods have been proposed^{14,15}. However, these approaches have mainly been designed for the Ising model and may not be ideal when handling realistic magnetic materials with long-range interactions.

Here, inspired by the idea of stochastic tunneling¹⁶ hybrid metaheuristic approaches^{17,18}, and evolutionary approaches^{14,15,19–21} a metaheuristic energy minimization approach is proposed and tested for magnetic systems with topological magnetic textures, with the goal of efficiently identifying the global minimum (a zero kelvin property). Specifically, it combines an evolutionary algorithm with a local optimizer. The evolutionary algorithm provides a way to select which new regions of the PES to explore, and the local optimizer finds the nearest local minima in the selected regions. We have selected to name our method genetic tunneling optimizer (GTO) to highlight its ability to tunnel energetically inaccessible regions¹⁶ of the PES using not only mutations but also cross-over operators i.e., a genetic approach. We analyze the performance of our proposed algorithm, which is designed to have a high ability to escape from local traps during energy minimization. In particular, we investigate the efficiency of our algorithm by performing simulations of a 2D monolayer with model exchange parameters that give rise to Bloch-type skyrmions as well as the experimentally well-studied Pd/Fe/Ir(111) system, which exhibits a Néel-type skyrmionic phase. Our results show that the proposed approach is indeed highly efficient in finding the global minimum in each of these test systems.

Results

Spin system parameterization. Searching for the ground state of a spin system at zero kelvin can be reformulated as finding the

global minimum of the PES defined by a many-body Hamiltonian describing the magnetic interactions between the constituents of the system. In the present work, we use a Heisenberg-type classical atomistic spin Hamiltonian of the form

$$\mathcal{H} = - \sum_{i \neq j} J_{ij} \mathbf{S}_i \cdot \mathbf{S}_j - \sum_{i \neq j} \mathbf{D}_{ij} \cdot (\mathbf{S}_i \times \mathbf{S}_j) - \sum_i \mathbf{B}^{\text{ext}} \cdot \mathbf{S}_i - \sum_i K^U (\mathbf{S}_i \cdot \mathbf{e}_z)^2, \quad (1)$$

where \mathbf{S}_i is the spin moment at site i . J_{ij} , \mathbf{D}_{ij} , K^U , \mathbf{e}_z , and \mathbf{B}^{ext} are Heisenberg exchange interactions, Dzyaloshinskii-Moriya interactions, uniaxial anisotropy, the easy axis vector, and the applied field, respectively. Typically, these four Hamiltonian terms are sufficient for a good description of the system one wishes to analyze (see, e.g., ref. ⁶). Note that for some magnetic systems, additional terms may be relevant to include, e.g., the biquadratic exchange coupling or the four-ring-interaction⁶.

In this work, we test our algorithm by studying two different systems—a model system and a system with realistic materials-specific magnetic interaction parameters. In the model system, we use only nearest-neighbor interactions, and the parameters are chosen to generate Bloch-type skyrmions. The second system is the Pd/Fe/Ir(111) system. Here, we included all interactions within the 30 nearest neighbor shells. All parameters were calculated by means of ab-initio DFT^{22,23}, and can be found in Supplementary Note 7.

In the present work, we have considered spin configurations of 100×100 atomic spins, but this number can naturally change, depending on the studied system. All interaction data for both systems can be found in the GitHub repository (See data availability section).

Outline of the genetic tunneling procedure. Finding the global minimum in a complex PES of a spin system such as the one described above is commonly a non-deterministic polynomial-time hard (NP-hard) problem¹³. These problems are very challenging, which has prompted the development of various heuristic methods, for example simulated annealing (SA). To solve this global energy optimization problem, we propose a genetic tunneling algorithm. The complete details are described in the Method section, here, we outline the most salient features. The procedure is illustrated schematically in Fig. 1, where, in particular, a flow chart of the method can be found in panel (d). Our approach builds on genetic algorithms that serve as an optimization scheme by mimicking the flow of genetic material through generations with an evolutionary tendency for finding better solutions. In our method, segments of spin configurations are conceptually represented as genes, each assigned a quality tag in terms of the underlying energy. In this way, following the logic of genetic algorithm-based optimisation, our method aims at finding the global minimum. Importantly, the overall aim of the genetic algorithm is to allow us to reach the global minimum with minimal numerical effort, while we avoid becoming trapped in metastable configurations (i.e., local minima).

The workflow used here can be summarized into two parts: finding spin configurations corresponding to local minima in the PES and using these spin configurations to perform genetic tunneling over the PES.

The very first step in the procedure is of course to provide basic information about the system to the algorithm—input which consists of physical information such as crystal lattice, system size, atomic positions, and magnetic interactions. An initial set of preliminary spin configurations is then created. These spin configurations serve as a first coarse guess and are typically far from any energy minimum. For each of these preliminary configurations, a local optimization module is invoked, relaxing

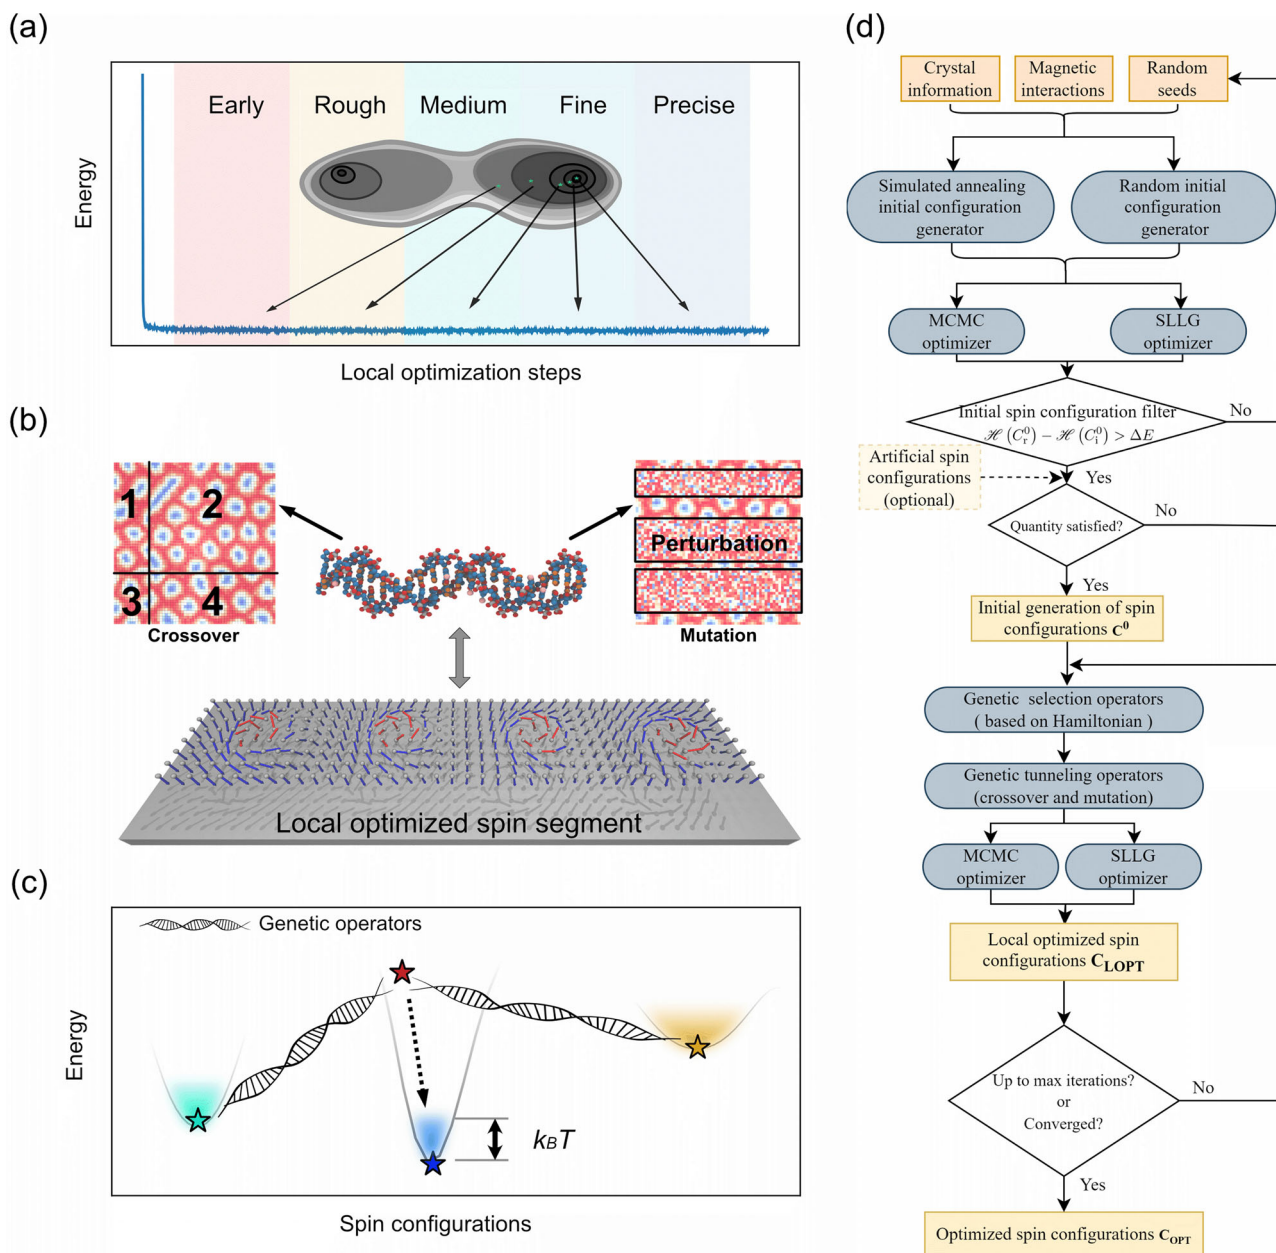


Fig. 1 Conceptual illustration of the genetic tunneling optimizer. **a** Conceptual illustration of the variance-threshold controlled localized optimization process to find low-energy spin configurations. The contour map in the middle shows an example of a potential energy surface (PES). The darker zones in the contour map represent lower energy areas. The arrows connect the local optimization process and the configuration point in the PES. At the top of the figure, five colored blocks, i.e., Early, Rough, Medium, Fine, and Precise, denote different converge levels of the local search algorithm. The Early level means the lowest convergence and the Precise level represents the highest convergence. **b** Conceptual illustration of how the genetic operators are applied to the spin system. The whole process involves three subprocesses, i.e., spin-configuration segmentation, crossover, and perturbation-based mutation (for details, see the Method section). In the spin configuration segmentation part, different from conventional binary-based genetic representation, the configuration-space spin textures are viewed as information carriers conceptually similar to chromosomes in biological systems. The spin textures are divided into several segments that can be used in the same way as gene segments. Examples of the square-crossover and mutation operators are shown at the top of Fig. (b) (for details about these genetic tunneling operators, see the Method section). **c** Conceptual illustration of how the genetic tunneling operators tunnel through the energy barriers and enable a heuristic search for spin configurations with lower energy, with the ultimate aim of finding the global minimum. The curve represents the PES, and the colored stars represent optimized spin configurations corresponding to local minima of the PES at temperature T . The height of the shadow-colored region equals $k_B T$. **d** Flowchart of the entire procedure. The dark yellow boxes at the top represent input data that need to be prepared before execution, and the light yellow boxes represent generated spin configurations. The light yellow box with a dashed boundary represents an optional choice. The gray rounded rectangles represent operations. The white diamond-shaped boxes represent conditional statements. Notation in the flow chart: $\mathcal{H}(C_i^0)$ and $\mathcal{H}(C_i^1)$ represent the energy of a generated spin configuration corresponding to a local minimum in the PES and the energy of any spin configuration that has already been selected as part of the initial parent generation, respectively. ΔE is a threshold energy difference set so that spin configurations that are too similar to the already selected ones are discarded—see Eq. (8) for more details. C^0 is the initial sets of selected spin configurations, constituting the first parent generation. Finally, C_{LOPT} and C_{OPT} are the sets of local and final optimized spin configurations, respectively. MCMC and SLLG optimizer represent Markov chain Monte Carlo and spin-dynamics-based local optimizers, respectively.

all magnetic orientations so as to reach, for each of them, the closest local minimum. A selection also takes place so that the spin configurations that become members of the first parent generation are not too similar to each other. The end result of this step is a diverse set of initial-generation spin configurations representing local minima in the PES. Then follows the metaheuristic search. Here, the set of initial-generation spin configurations representing local minima is segmented into pieces. The pieces are then combined together in new ways and also subjected to some mutation by adding random noise, see Fig. 1b. This allows us to reach new, unexplored parts of the PES. The process is then repeated until convergence, i.e., new local minima are identified, segments of the corresponding spin configurations are combined, and yet new parts of the PES are reached, with the aim of identifying spin configurations with still lower energies than the ones in the previous generations. Additional details on the various parts making up the genetic tunneling method can be found in the Method section.

Algorithm performance comparisons and optimization. Each different part of the method proposed here can be implemented in various ways. In order to optimize the performance of our method we have therefore systematically combined a number of possible implementations of each part and compared the resulting performance. Specifically, we combine three typical genetic selection operators, i.e., Rank, Tournament, and Roulette Wheel, with two different local optimizer backends – the MCMC-based optimizer and an optimizer based on the stochastic Landau-Lifshitz-Gilbert (SLLG) equation. Finally, for the purpose of analyzing the impact of starting from a pre-optimized initial spin configuration, an initial spin configuration generated with SA was tested as an initial guess, in addition to simply using a randomly generated initial spin configuration.

We have set the temperature to a very low value, 0.1 mK, in order to obtain highly converged energies and spin configurations of the local minima. This enables a more fair comparison of the performance of the various implementations, since random noise in the optimization process is reduced. For an explanation of the role of temperature in the algorithms and how we use the term ground state in relation to that, see the Method section. The simulations are performed for a spin Hamiltonian describing the 2D Pd/Fe/Ir(111) system.

The results are presented in Fig. 2. To facilitate comparison, a classical MCMC-based SA optimization with a fine temperature mesh and 2.5×10^6 steps is used as baseline. These baseline results are denoted SA in Fig. 2a, and the SA predicted spin configuration is shown in Fig. 2c. For more details on the SA setup we have used, see the Method section. A clear trend in Fig. 2a is that all tested combinations indeed manage to find spin configurations with gradually lower energy with each iteration. However, there is a significant difference in performance between the MCMC-based implementations and the spin-dynamics-based ones. Specifically, we find that with a finite number of local optimization steps and global searching epochs, all spin-dynamics-based optimizers may not achieve convergence. In general, the implementations with an MCMC-based backend perform significantly better, yielding consistently lower energies for the same number of iterations. In fact, all tested genetic tunneling operators combined with the MCMC backend invariably reach convergence and find solutions with lower energy compared to the other tested methods, including the baseline reference energy obtained by SA. We also measured and compared the computing time for all implementations, and found that the ones with the MCMC backend were faster on average, see Fig. 2b.

Figure 2d–k shows the spin configurations after 60 iterations for each tested implementation. All implementations with the MCMC backend (Fig. 2d–g) find a hexagonal skyrmion lattice with a unit cell size in good agreement with previous studies^{24,25}. In contrast, none of the implementations based on the spin-dynamics backend (Fig. 2h–k) manage to identify a stable hexagonal skyrmion lattice within 60 iterations. The better efficiency of the Monte Carlo-based simulations are natural since, in the form employed here, they are designed and used to find minimum configuration of atomic moments, while the spin-dynamics equations are designed to follow the real-time equations of motion of a physical object, in this case an atomic spin. Hence, the latter attempts to follow the time evolution of an object, in many cases not even attempting to find a ground state, while the former are tailored to quickly find the lowest energy state.

All in all, we find that the Rank selection genetic operator algorithm in combination with the MCMC backend appears to be the best choice. Furthermore, we find that starting from a completely random spin configuration is a good choice, since no significant gain could be identified in the tests where we instead started from an initial configuration found using SA. Thus, in the remainder of this work, we use the algorithm corresponding to option (d) in Fig. 2, i.e., MCMC with Rank selection and random initial input, and refer to it as GTO-MCMC-Rank.

Effect of system size and applied field on performance. In Fig. 3, we summarize how our method performs as a function of system size and applied magnetic field, and compare with the corresponding performance of SA. In all these simulations the temperature was set to a low value of -0.1 mK and the tests were done for the Fe/Pd/Ir(111) system. Fig. 3a shows total energy (left y-scale) and execution time (right y-scale) as a function of system size. Here, the magnetic field was set to 2.7 T, and all simulations were performed on a single node with 32 CPU cores.

We find that GTO-MCMC-Rank consistently identifies lower energy states compared to SA for all tested system sizes. GTO-MCMC-Rank converges faster than SA for system sizes up to about 160×160 but scales less well than SA for larger system sizes. We also evaluated the performance of the GTO-MCMC-Rank as a function of applied field, see Fig. 3b. The left y-scale shows the total energy of the system, and the right y-scale the total topological charge (or skyrmion number). Just as before, we see that GTO-MCMC-Rank consistently identifies lower energy states compared to SA.

The most remarkable result here is however how GTO-MCMC-Rank very successfully can be used to distinctly identify the three different phases of the system and their regions of stability with respect to the magnetic field. This is because the method predicts the topological charge with excellent precision. Below 1.5 T the spin-spiral state is the ground state and above 3.3 T the ground state is ferromagnetic, both with zero topological charge. From about 1.6 T to around 3 T, a hexagonal skyrmion lattice phase is predicted, with a constant topological charge of 36 (corresponding to a 6×6 skyrmion lattice) for the investigated system size (the process of predicting under 0.7 T, 2.7 T, and 3.7 T can be found in supplementary Movie 1, supplementary Movie 2, and supplementary Movie 3, respectively.). In contrast, the SA simulations predict nonzero topological charge over the entire investigated magnetic field interval. Thus, SA neither identifies the three different phases, nor the intermediate regions. Moreover, the topological charge computed with SA oscillates visibly in the region where there should be a stable hexagonal skyrmion lattice. In the regions where the topological charge should be zero, SA predicts a topological charge only slightly different compared to the skyrmion-lattice region.

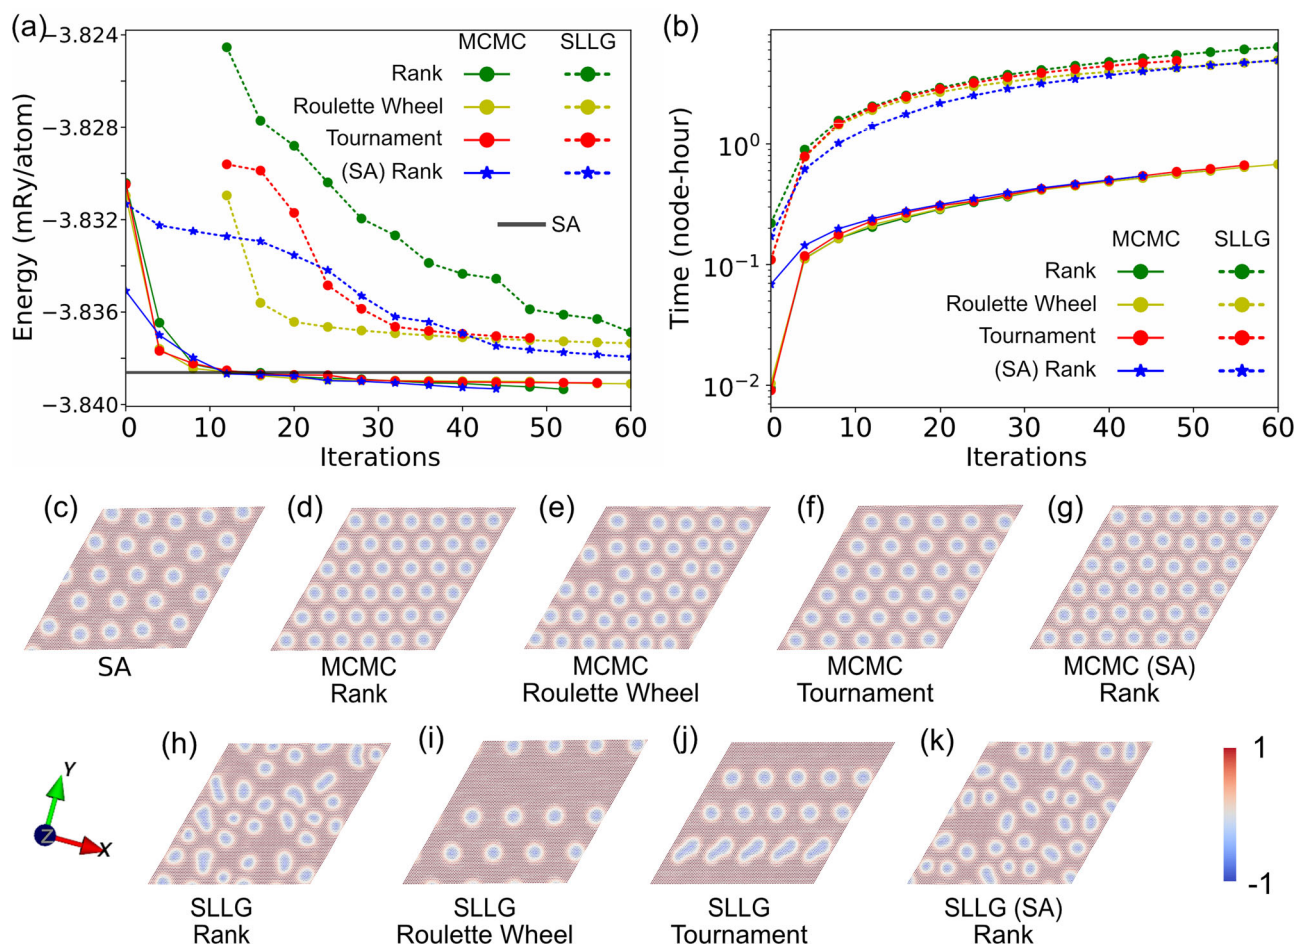


Fig. 2 Performance analysis of the genetic tunneling optimizer with different local optimizers. **a** Performance benchmark of the genetic tunneling optimizer at 0.1 mK. The simulation system is Pd/Fe/Ir(111), with an external magnetic field of 2.7 T directed out of the plane. Here Markov chain Monte Carlo (MCMC) and SLLG represent Markov chain Monte Carlo and spin-dynamics-based local optimizers, respectively. In this work, the SLLG local optimizer solves the stochastic Landau-Lifshitz-Gilbert (SLLG) equation with an artificial damping value of 0.4. Rank, Roulette Wheel, and Tournament represent genetic selection methods (for details, see the Method section). The abbreviation SA stands for simulated annealing, the results of which are used as reference energy and only used in the first generation in this work (See Supplementary Note 4). For better comparison, the results from SLLG with Rank, Roulette Wheel, and Tournament selection methods are shown from generation 12 of the algorithm. **b** The time consumption in units of node-hours for each optimizer. All of the simulations are performed on an Intel Xeon Gold 6130 CPU node with 32 cores without concurrent processes. **c–k** Visualization of the final spin configurations from each optimizer. Here, blue color means that the spin points downwards away from the reader, whereas red color means that the spin points upward, toward the reader. The white color indicates that the spin direction is parallel to the plane.

The GTO-MCMC-Rank simulations developed here predict that the zone where the skyrmion lattice is the ground state is somewhat wider than what was found in ref. 22. The blue and gray areas indicate transition zones in the GTO-MCMC-Rank simulations, in which the topological charge changes as the system changes its ground state. Energy differences between configurations with various topological charges in the second transition zone can be found in Supplementary Note 2, 6. Finally, in Fig. 3c, d, we show how the GTO-MCMC-Rank execution time scales with the number of CPU cores used. As already mentioned, our method lends itself well to concurrent computing, which becomes increasingly important the larger the addressed system is. In both figures, the left y-scale is execution time and the right y-scale is the computed total energy of the system. We show user execution time (the dark green circles) as well as core execution time, i.e., the user execution time multiplied with the used number of cores (dark blue circles). In Fig. 3c, the size of the tested system is 100×100 , whereas in Fig. 3d it is four times larger— 200×200 . As regards the computed total energy per atom (the light green circles), it is found to be effectively constant, as

expected. From the core execution time data, we see that the method scales very well with the number of cores. For the 100×100 system in Fig. 3c, the core execution time is smallest for 64 cores, but user execution time can still be gained up to 256 cores. For the 200×200 system, the corresponding numbers are 256 and 512.

Computed ground state as a function of applied magnetic field with constant temperature.

In this section, we present results from investigations on how the genetic tunneling algorithm performs at a temperature of 8 K for different applied fields. Results for two simulated systems are shown in Fig. 4. The left column (Fig. 4a–c), shows simulations on an artificial frustrated spin system that exhibits a spin spiral state (a), a Bloch-type skyrmion state (c), and ferromagnetic state (e) at low, medium, and high applied field, respectively. The simulated system in the right column in Fig. 4b, d, f is a Pd/Fe/Ir(111) monolayer, which also has three states at different fields, but instead of Bloch skyrmions, it contains Néel-type skyrmions.

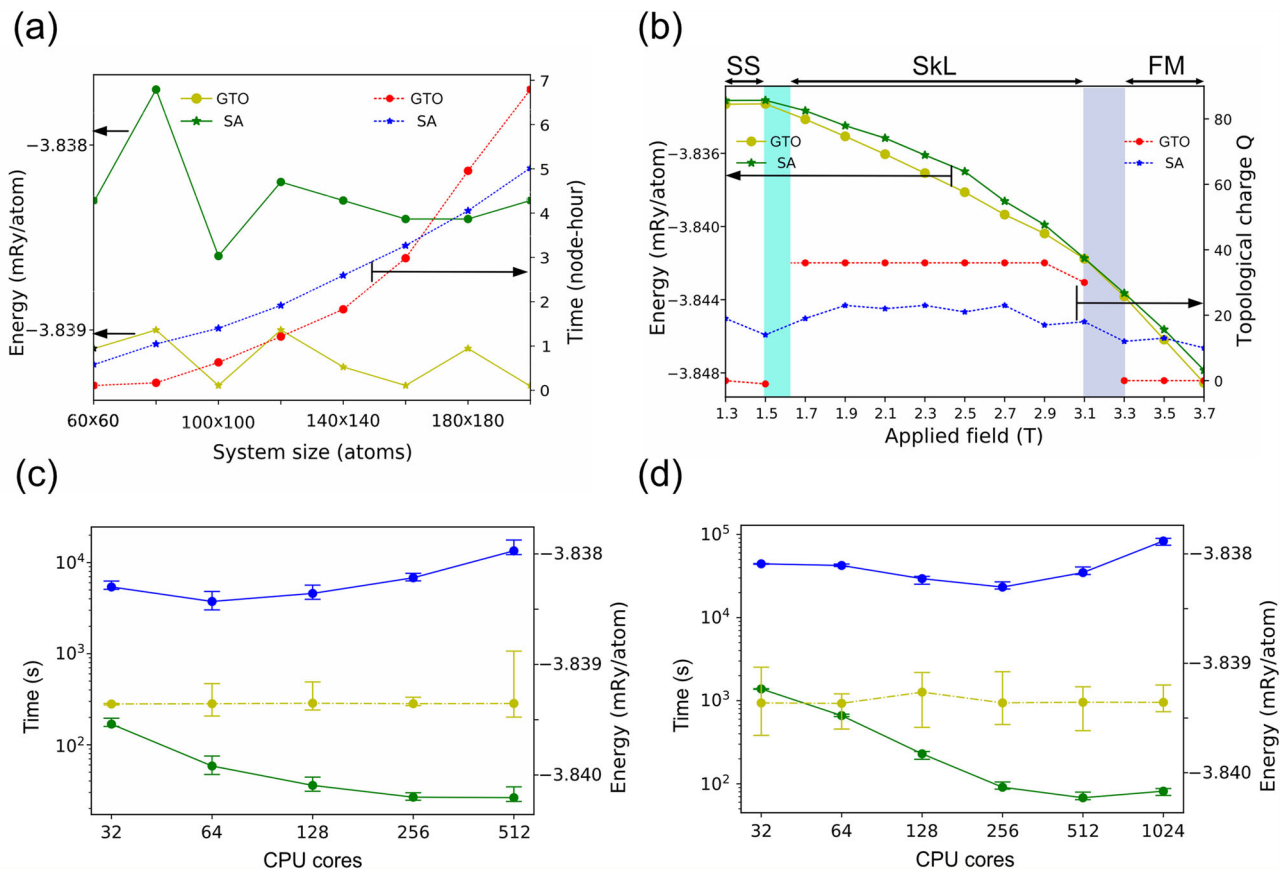


Fig. 3 Performance analysis of the genetic tunneling algorithm at low temperature on the Fe/Pd/Ir (111) system with different system sizes and applied fields. **a** Predicted ground-state energy and simulation execution time as a function of system size. Due to limited space in the figures, we use the label GTO to refer to GTO-MCMC-Rank computations. SA stands for simulated annealing. Here MCMC and GTO represent Markov chain Monte Carlo-based local optimizers and the genetic tunneling optimizer, respectively. **b** Total energy and topological charge as a function of applied magnetic field. The size of the simulated system is 100×100 , and the temperature is set to 0.1 mK. SS, SkL, and FM represent the spin spiral, skyrmion lattice, and ferromagnetic states, respectively. (Discussion on system size, see Supplementary Note 5) The cyan- and purple-colored blocks indicate transition zones between the states. **c** First-generation simulation user execution time (green symbols) and core execution time (blue symbols), and predicted ground-state energy (yellow symbols) as a function of the number of CPU cores for a 100×100 spin system. The lines are guides for the eye. **d** Same as in (c), but for a 200×200 spin system. Each point in (c) and (d) represents five simulations. The time error bars in (c) and (d) show the highest, lowest, and average first-generation execution time. The energy error bars in (c) and (d) indicate the highest, lowest, and average predicted ground state energies.

As shown in Fig. 4a–c, both the GTO-MCMC-Rank and SA can find the ground state of the artificial system with different applied fields, but as regards the energy itself, the genetic tunneling algorithm is more successful in finding spin configurations with lower energy. This result is also reflected in the real-space spin configuration visualization, especially in Fig. 4b. Evidently, while magnetic moments in both systems follow the Boltzmann distribution at 8 K, the hexagonal skyrmion lattice found with the genetic tunneling algorithm is more stable.

From Fig. 4d–f, we find that also for a more complex, real system, the GTO-MCMC-Rank still performs better than SA. When the applied field is 1 T, the ground state predicted by the GTO is a twisty spin spiral state which is in remarkably good agreement with the experimental data in ref. 25. However, SA gives a hybrid phase that contains both spin spirals and bubbles. At 2.7 T, GTO-MCMC-Rank finds a more reasonable hexagonal skyrmion lattice compared to SA. Finally, at 4.4 T, the GTO-MCMC-Rank correctly finds the ferromagnetic state (with some fluctuations due to finite temperature) whereas in the SA solution there are several complex structures (bubbles or skyrmions). To summarize, the final system energy that we obtain from the GTO is lower than what we obtain using SA in both systems addressed.

Thus, we conclude that our algorithm works well for a range of magnetic fields in the presence of thermal fluctuations.

Computed ground state as a function of temperature with constant applied magnetic field.

Next, we analyzed the performance of our algorithm as a function of temperature. In all these simulations, the applied magnetic field was set to 2.7 T (i.e., in the middle of the interval where the skyrmion lattice is the ground state) and the tests were done for the Fe/Pd/Ir(111) system. The results are shown in Fig. 5. In this case, we find that the ground state energies predicted by GTO-MCMC-Rank (yellow symbols) and SA (green symbols) are broadly similar. However, as the temperature increases, the difference increases, with GTO-MCMC-Rank consistently finding slightly lower energy states, see Fig. 5a. We also computed the topological charge of the predicted ground states. In the GTO-MCMC-Rank simulations, the topological charge decreases with temperature, which is consistent with experimental results^{25,26} and also what one would expect. In contrast, SA does not reproduce this trend. Additional data in the form of spin configurations are shown in Fig. 5b–i, where the top row is the GTO-MCMC-Rank spin configuration results, and the

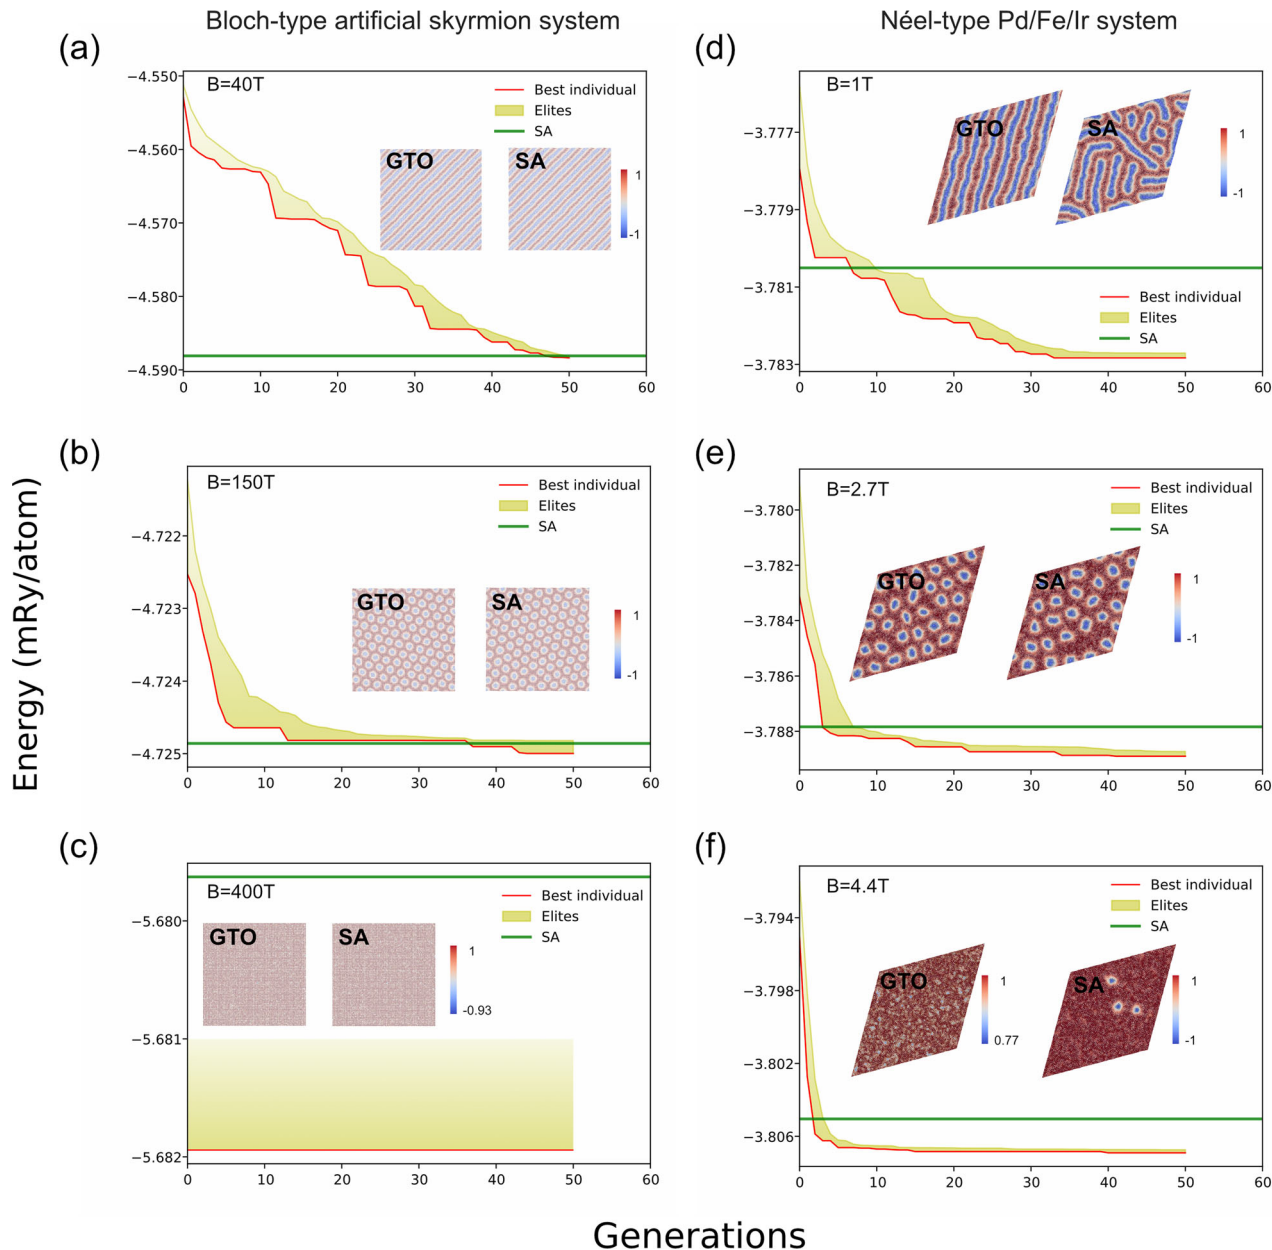


Fig. 4 Comparison of ground state searching between simulated annealing and genetic tunneling optimizer. Searches for the ground state in an artificial Bloch-type skyrmion system and a Néel-type skyrmion system Pd/Fe/Ir(111), using a simulation temperature of 8 K. Due to limited space in the figures, we use the label GTO (genetic tunneling optimizer) to refer to GTO-MCMC-Rank computations where MCMC represents Markov chain Monte Carlo based local optimizers. SA stands for simulated annealing. **a-c** These show the simulation results for the artificial Bloch-type skyrmion system with 40 T, 150 T, and 400 T applied fields, respectively. **d-f** These show the simulation results for the Pd/Fe/Ir(111) system with 1.0 T, 2.7 T, and 4.4 T applied field, respectively. The ground states in **(a-d)**, **(b-e)**, and **(c-f)** are a spin spiral, a skyrmion lattice, and the ferromagnetic state, respectively. In each panel, the red line represents the energy of the best individual of each generation, the yellow band represents the energy distribution of elite individuals^{34,35} (See Supplementary Note 3) in each generation, and the green line represents the energy of the spin configuration predicted using SA. To show the optimization process in detail, we set the convergence limit to an extremely low value. Based on this, all optimizations will run up to the preset maximum iteration threshold of 50 in this section. The real-space spin configurations found by GTO-MCMC-Rank and SA are visualized in the middle of each panel. Here, blue color means that the spin points downwards away from the reader, whereas red color means that the spin points upward toward the reader. The white color indicates that the spin direction is parallel to the plane.

bottom row is the corresponding SA results. In the GTO-MCMC-Rank images, remnants of the ground-state hexagonal skyrmion crystal pattern can be seen up to 32 K, whereas such patterns are absent in the SA results above 8 K. It is reasonable to expect that the zero kelvin hexagonal skyrmion lattice breaks up gradually with increasing temperature, and therefore, we believe that these images further strengthen the conclusion that GTO-MCMC-Rank has superior performance compared to SA.

Discussion

To the best of our knowledge, this is the first study establishing a genetic-tunneling optimization protocol for complex spin systems with long-range interactions. The algorithm presented in this work has general applicability to spin systems, and is here shown to successfully find the ground state for monolayer spin systems at finite temperatures and at a variety of applied magnetic fields. The approach contains two essential parts: (1) a

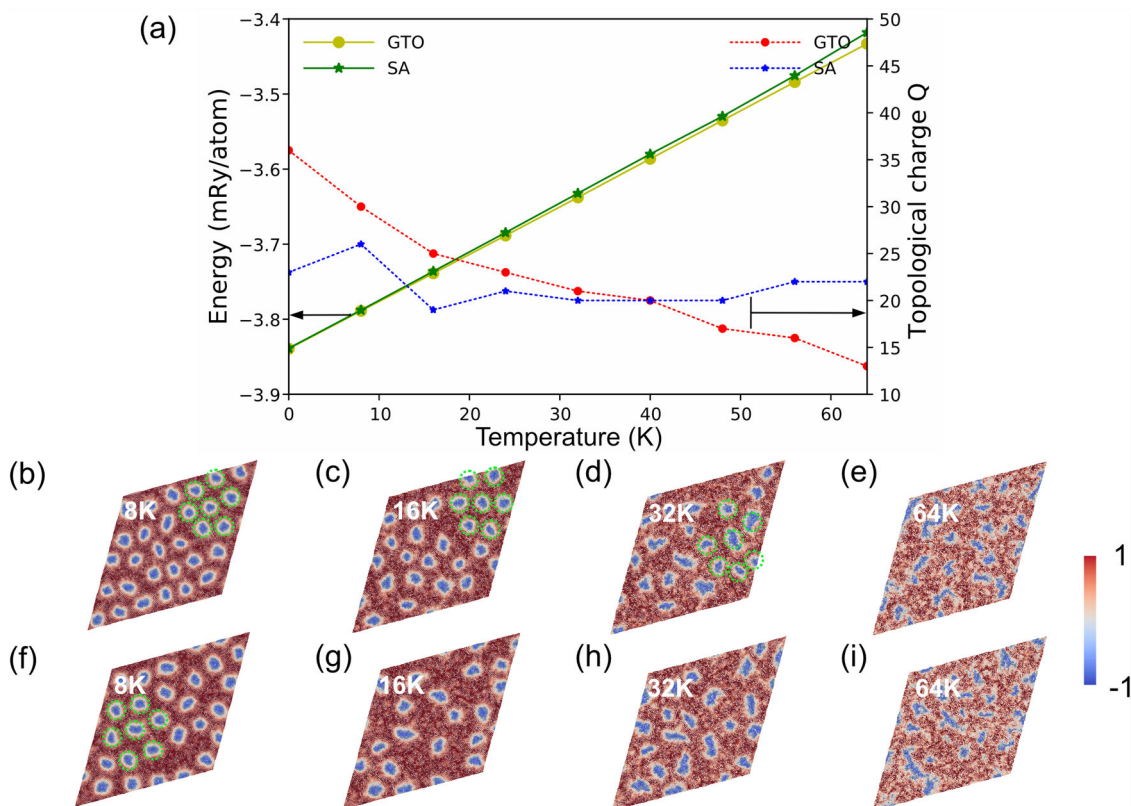


Fig. 5 Ground-state search for the Pd/Fe/Ir (111) system as a function of temperature. **a** The ground state energy (left y-axis) and topological charge (right y-axis). The applied magnetic field is set to 2.7 T. Due to limited space in the figures, we use the label GTO to refer to GTO-MCMC-Rank computations. Here MCMC and GTO represent Markov chain Monte Carlo-based local optimizers and the genetic tunneling optimizer, respectively. SA stands for simulated annealing. **b–e** Real-space plots of ground-state spin configurations identified using GTO-MCMC-Rank for four different temperatures. **f–i** Real-space plots of ground-state spin configurations identified using SA for four different temperatures. The green dashed lines point out patterns with hexagonal ordering of skyrmions. The color scheme in the plots (**b–e**) is the same as in Fig. 2c–k, where blue color means that the spin points downwards away from the reader, whereas red color means that the spin points upward, toward the reader. The white color indicates that the spin direction is parallel to the plane.

variance-threshold controlled local optimizer, which includes an MCMC optimizer and a spin-dynamic optimizer, and (2) a spin-configuration space genetic tunneling metaheuristic search module. The algorithm is designed to be able to escape from a local minimum through the use of genetic tunneling operators and find the global minimum for a given system without any initial guess. The efficiency of our genetic tunneling protocol is investigated on both a simple artificial system with magnetic frustration and a Pd/Fe/Ir(111) monolayer that includes complex Heisenberg and Dzyaloshinskii-Moriya interactions, as calculated from Density functional theory (DFT). The results indicate that the GTO-MCMC-Rank has better performance than SA when it comes to finding stable spin configurations as a function of external parameters like temperature and applied magnetic field. Most noteworthy, we here considered the spiral structure, the skyrmion lattice, and the ferromagnetic state. Our method successfully finds all three ground states as a function of magnetic field, and also correctly identifies the transition regions in between. To our knowledge, no other theoretical method has yet demonstrated such an ability. It can also be concluded that the performance of the algorithm is limited neither by the system size, geometry, nature of the magnetic interactions, temperature, nor the applied field strength.

In practice, for optimal performance of the protocol, the various hyperparameters will of course need to be fine-tuned depending on the system under consideration. For example, a non-optimal variance threshold may increase the risk of premature convergence.

In conclusion, we have explored a genetic tunneling protocol, which is designed to predict the magnetic ground state of a classical spin Hamiltonian at finite temperature. We demonstrate that our method is robust for 2D systems, both for simpler model systems and for the more complex Pd/Fe/Ir(111) system. We envision that our findings will pave the way for evolutionary computing in finding the ground state of magnetic systems, e.g., for magnets with non-trivial topology and spin glass systems. Since complex systems in a very general sense can be cast into the language of spin Hamiltonians, it also appears possible that the here suggested protocols will find applications in other areas of solid-state science, or even in fields outside the natural sciences.

Methods

A metaheuristic algorithm can interactively guide and modify the operations of subordinate heuristics, to efficiently produce preferable solutions within a high-dimensional search space^{27,28}. Representative metaheuristic algorithms include SA, Particle Swarm Optimization²⁹, and Genetic Algorithms (GA)³⁰. It has been shown that a hybrid algorithm that combines the methods mentioned above with a local optimizer, e.g., gradient descent, can be an efficient way to solve the global optimization problem within a complex configuration space³¹.

In this study, we introduce a genetic tunneling strategy in the form of a hybrid algorithm—connecting the local minimization approach shown in Fig. 1a to a metaheuristic genetic-tunneling module—with the aim of finding the global energy minimum of

the spin system. For an overview of the procedure, see the flow chart in Fig. 1d. In the following sections, we describe the main implemented algorithms and how they interact.

Selection and genetic-tunneling operators. Below, we describe the evolutionary parts of our hybrid algorithm. Each generation consists of N_p spin configurations, i.e., for the k th generation we have

$$\mathbf{C}^k = \{C_i^k | i = 1, 2, \dots, N_p\}, \quad (2)$$

where \mathbf{C}^k is the set of all spin configurations in generation k , and C_i^k is one spin configuration. Starting from the k th generation of N_p spin configurations, each corresponding to a local minimum in the PES, generation $k + 1$ is produced in the following way. The first step is to select spin configurations in generation k that will be used as parents for a new spin configuration in generation $k + 1$. We have implemented three different selection operators—Roulette Wheel, Rank, and Tournament. With the Roulette Wheel selection operator, the probability of choosing a candidate spin configuration C_i^k to become a member of the breeding pool is equal to

$$P(C_i^k) = \frac{\mathcal{H}_i}{\sum_{i=1}^{N_p} \mathcal{H}_i}, \quad (3)$$

where \mathcal{H}_i is the Hamiltonian (i.e., the energy) of spin configuration C_i^k . (Note that both the numerator and the denominator on the right-hand side in Eq. (3) are negative.) The Rank and Tournament selection operators are described as algorithms in Supplementary Note 1, where a graphic indicator demo of the Roulette Wheel can also be found. In the present implementation, N_p is set to 64, and we select a total of four spin configurations out of these to breed one new spin configuration.

The next step is to create the new spin configuration from the selected parents. To perform this step, we invoke three tunneling operators—square crossover, linear crossover, and mutation. The pseudocode for the squared-crossover and linear-crossover operators can be found in Supplementary Note 1. In the square-crossover scheme, we first generate two random split ratios and use them to split each parent spin configuration along the x - and y -directions into four rectangles, see Fig. 1b. Then, we combine these generated parts using a randomized procedure to create a new spin configuration with the same dimensions as the spin configurations in the parent generation. Finally, the mutation operator is added. This operator consists of adding a Gaussian random noise to all spins in the new configuration, using the Rodrigues rotation formula³². Thus, the square-crossover scheme consists of two operators—the square-crossover operator and the mutation operator. The linear-crossover scheme is very similar to the square-crossover scheme, with the difference that we instead split each parent spin configuration along only one direction. These steps are then repeated until we have N_p new spin configurations. As a side note, we mention that the square-crossover and linear-crossover tunneling operators we introduce here are especially well-suited for systems with clearly contained spin textures, such as skyrmions and spin spirals.

Before the new generation of spin configurations can be used as parents, a local energy optimizer is invoked to guarantee that all spin configurations in the new generation correspond to local minima in the PES. The local energy optimizers we have used for this step are described next.

Variance-controlled local energy optimizer at finite temperature. A local optimization of a spin configuration involves finding a local minimum starting from a given initial guess. We have

implemented two different types of local optimizers: an MCMC-based optimizer and one based on SLLG. These two methods have proven robust and efficient in describing complex spin systems^{6,33}. We describe both in some detail below.

The MCMC optimizer is of Metropolis type and performs energy minimization under finite temperature by using the transition probability P_t between two spin configurations in the Markov chain:

$$P_t = \begin{cases} \exp\left(-\frac{\Delta E}{k_B T}\right), & \text{if } \Delta E > 0 \\ 1, & \text{otherwise} \end{cases} \quad (4)$$

where ΔE , k_B , and T are the energy difference between spin configurations, the Boltzmann constant, and the temperature of the system. For a given initial spin configuration, the method will iteratively minimize the energy of the system.

The second approach we use to optimize spin configurations is the spin-dynamics-based optimizer, which uses the SLLG equation to simulate the time evolution of atomic magnetic moments. This method is able to reach a spin configuration near a local energy minimum from a given initial state since—when Gilbert damping³³ is included in the simulations—energy is allowed to dissipate from the system. The atomistic SLLG equation reads

$$\frac{d\mathbf{m}_i}{dt} = -\gamma_L \mathbf{m}_i \times (\mathbf{B}_i + \mathbf{B}_i^f) - \gamma_L \frac{\alpha}{m_i} \mathbf{m}_i \times [\mathbf{m}_i \times (\mathbf{B}_i + \mathbf{B}_i^f)], \quad (5)$$

where \mathbf{B}_i the effective magnetic field at site i and \mathbf{B}_i^f is the stochastic magnetic field corresponding to the thermal fluctuations present in a heat bath with temperature T . In this equation, the first term represents the precessional motion of the atomic magnetic moments, while the second term describes the damping motion. In the expression above, γ_L is the renormalized gyromagnetic ratio, which is calculated from

$$\gamma_L = \frac{\gamma}{(1 + \alpha^2)}, \quad (6)$$

where γ is the gyromagnetic ratio and α is the isotropic Gilbert damping constant. The performance of the SLLG local optimizer is highly related to the value of damping and time step. In the present work, we use both the Metropolis MCMC and the SLLG local optimization schemes as implemented in the Uppsala Atomistic Spin Dynamics (UppASD) package¹⁰.

In order to automatically stop the minimization process when a desirable convergence level has been reached, we use a variance threshold $\text{Var}(\mathcal{H})$. For the spin Hamiltonian used here, this threshold is defined as

$$\text{Var}(\mathcal{H}) = \frac{1}{n} \sum_{i=1}^n (\mathcal{H}_i - \langle \mathcal{H} \rangle)^2, \quad (7)$$

where $\langle \mathcal{H} \rangle$ stands for the expectation value of the spin Hamiltonian, and the sum is over the n last iteration steps. In the present work, we typically set n between 10 and 100.

We end this subsection with a note on nomenclature: as is evident from the above, in our simulation method the temperature must always be set to a finite value, it cannot be zero. Therefore, in this work, the term ground state should be taken to mean the lowest lying minimum at the set simulation temperature, see Fig. 1c.

Initialization and termination

Initialization. The algorithm starts from an initial parent generation \mathbf{C}^0 of spin configurations (generation zero). We have implemented two ways of producing these initial configurations. The first method starts by simply producing random spin

configurations using random numbers together with the crystal information and magnetic interactions defining the studied spin system. The generated spin configurations are subsequently relaxed to the closest lying local minimum using one of the local energy optimizers described above. We select N_p spin configurations by using the criterion that any two configurations must not be too close in energy, i.e.,

$$\mathcal{H}(C_r^0) - \mathcal{H}(C_i^0) > \Delta E. \quad (8)$$

Here, ΔE is a threshold energy difference that guarantees that configurations C_r^0 and C_i^0 are not too similar. In this work, we use values ranging from 10^{-4} to 10^{-6} mRy/atom for ΔE . In practice, this implies that the number of random spin configurations generated will typically be much larger than N_p , in order to be able to select N_p spin configurations fulfilling the selection criterion above. The N_p selected spin configurations constitute \mathbf{C}^0 . The second method is based on SA. Here, we use a relatively coarse temperature mesh so that the SA-based initialization does not become too time consuming. Just as in the first method, a large number of spin configurations are generated, a local minimizer is invoked to find the corresponding local minima, and N_p spin configurations are selected using the criterion in Eq. (8). Specifically, in the present work, the SA simulations were performed using a temperature mesh of only four points (the simulated temperature, and then 20 K, 50 K, and 200 K added to the simulated temperature). At each temperature, we performed 2000 Metropolis steps. It is worth noting that here, the SA procedure is only used to generate the initial generation in the GTO optimization process. It is simply an alternative way to produce the first parent generation. From the second generation, no temperature annealing is performed.

Termination. There are two stop criteria set for the search process in this work—a maximum number of allowed iterations (i.e., number of produced generations) and a convergence criterion. In this work we have set the maximum number of iterations to around 60. However, our optimization algorithm often finds a converged solution (according to the convergence criterion explained below) within a significantly smaller number of iterations. Our convergence criterion is designed in the following way. For each generation, the variance over the set of spin configurations is computed. Thus, for each generation, we compute

$$\text{Var}(\mathcal{H}) = \frac{1}{N_p} \sum_{i=1}^{N_p} (\mathcal{H}_i - \langle \mathcal{H} \rangle)^2, \quad (9)$$

where $\langle \mathcal{H} \rangle$ stands for the expectation value of the spin Hamiltonians over all spin configurations in the set. As the optimization proceeds, the spin configurations within each generation will become more and more similar to each other as they approach the ground state. Therefore, when the variance in Eq. (9) decreases below a predefined threshold, the procedure is deemed to have converged and is stopped. At this point, the final spin configuration or configurations \mathbf{C}_{OPT} representing the best solution can be extracted (see the flowchart in Fig. 1d).

Computation of the topological charge. The topological charge Q of a 2D system in the xy -plane is computed using

$$Q = \frac{1}{4\pi} \int_A \vec{n} \cdot \left(\frac{\partial \vec{n}}{\partial x} \times \frac{\partial \vec{n}}{\partial y} \right) dx dy. \quad (10)$$

Here, $\vec{n}(x, y)$ is the unit vector of the local magnetization. The integral is taken over the area A in the xy -plane.

Data availability

All data needed for reproducing the results can be found in the GitHub repository <https://github.com/MXJK851/GTO-2D>, and Supplementary Movies 1, 2, and 3 are available within the paper.

Code availability

All code of GTO-2D is available at <https://github.com/MXJK851/GTO-2D> under the GPL-3.0 license.

Received: 12 April 2023; Accepted: 24 August 2023;

Published online: 02 September 2023

References

- De las Cuevas, G. & Cubitt, T. S. Simple universal models capture all classical spin physics. *Science* **351**, 1180–1183 (2016).
- Mühlbauer, S. et al. Skyrmion lattice in a chiral magnet. *Science* **323**, 915–919 (2009).
- Yu, X. Z. et al. Transformation between meron and skyrmion topological spin textures in a chiral magnet. *Nature* **564**, 95–98 (2018).
- Tchernyshyov, O. & Chern, G.-W. Fractional vortices and composite domain walls in flat nanomagnets. *Phys. Rev. Lett.* **95**, 197204 (2005).
- Speight, M. & Winyard, T. Skyrmions and spin waves in frustrated ferromagnets at low applied magnetic field. *Phys. Rev. B* **101**, 134420 (2020).
- Eriksson, O., Bergman, A., Bergqvist, L. & Hellsvik, J. *Atomistic spin dynamics: foundations and applications* (Oxford university press, 2017).
- Müller, G. P. et al. Spirit: multifunctional framework for atomistic spin simulations. *Phys. Rev. B* **99**, 224414 (2019).
- Hastings, W. K. *Monte Carlo sampling methods using Markov chains and their applications* (Oxford University Press, 1970).
- Roy, V. Convergence diagnostics for markov chain monte carlo. *Annu. Rev. Stat. Appl.* **7**, 387–412 (2020).
- Skubic, B., Hellsvik, J., Nordström, L. & Eriksson, O. A method for atomistic spin dynamics simulations: implementation and examples. *J. Phys. Condens. Matter* **20**, 315203 (2008).
- Verlhac, B. et al. Thermally induced magnetic order from glassiness in elemental neodymium. *Nat. Phys.* **18**, 905–911 (2022).
- Wang, W., Machta, J. & Katzgraber, H. G. Comparing monte carlo methods for finding ground states of ising spin glasses: population annealing, simulated annealing, and parallel tempering. *Phys. Rev. E* **92**, 013303 (2015).
- Hibat-Allah, M., Inack, E. M., Wiersema, R., Melko, R. G. & Carrasquilla, J. Variational neural annealing. *Nat. Mach. Intell.* **3**, 952–961 (2021).
- Chen, A., Choo, K., Astrakhantsev, N. & Neupert, T. Neural network evolution strategy for solving quantum sign structures. *Phys. Rev. Res.* **4**, L022026 (2022).
- Whitelam, S. & Tamblyn, I. Neuroevolutionary learning of particles and protocols for self-assembly. *Phys. Rev. Lett.* **127**, 018003 (2021).
- Wenzel, W. & Hamacher, K. Stochastic tunneling approach for global minimization of complex potential energy landscapes. *Phys. Rev. Lett.* **82**, 3003 (1999).
- D'Angelo, G. & Palmieri, F. Gga: a modified genetic algorithm with gradient-based local search for solving constrained optimization problems. *Inform. Sci.* **547**, 136–162 (2021).
- Kapoor, A., Nukala, E. & Chandra, R. Bayesian neuroevolution using distributed swarm optimization and tempered mcmc. *Appl. Soft Comput.* **129**, 109528 (2022).
- Hart, G. L., Blum, V., Walorski, M. J. & Zunger, A. Evolutionary approach for determining first-principles hamiltonians. *Nat. Mater.* **4**, 391–394 (2005).
- Zheng, F. & Zhang, P. Maggene: a genetic evolution program for magnetic structure prediction. *Comput. Phys. Commun.* **259**, 107659 (2021).
- Rahmanian Koshkaki, S. et al. Computational prediction of new magnetic materials. *J. Chem. Phys.* **157**, 124704 (2022).
- Miranda, I. P., Klautau, A. B., Bergman, A. & Petrilli, H. M. Band filling effects on the emergence of magnetic skyrmions: Pd/Fe and Pd/Co bilayers on Ir(111). *Phys. Rev. B* **105**, 224413 (2022).
- Bessarab, P. F. et al. Lifetime of racetrack skyrmions. *Sci. Rep.* **8**, 3433 (2018).
- Dupé, B., Hoffmann, M., Paillard, C. & Heinze, S. Tailoring magnetic skyrmions in ultra-thin transition metal films. *Nat. Commun.* **5**, 1–6 (2014).
- Romming, N. et al. Writing and deleting single magnetic skyrmions. *Science* **341**, 636–639 (2013).
- Lindner, P. et al. Temperature and magnetic field dependent behavior of atomic-scale skyrmions in Pd/Fe/Ir(111) nanoislands. *Phys. Rev. B* **101**, 214445 (2020).

27. Beheshti, Z. & Shamsuddin, S. M. H. A review of population-based metaheuristic algorithms. *Int. J. Adv. Soft Comput. Appl.* **5**, 1–35 (2013).
28. Abdel-Basset, M., Abdel-Fatah, L. & Sangaiah, A. K. Metaheuristic algorithms: A comprehensive review. *Computational Intelligence for Multimedia Big Data on the Cloud with Engineering Applications* 185–231 (Elsevier, 2018).
29. Bautu, A., Bautu, E. & Luchian, H. Particle swarm optimization hybrids for searching ground states of Ising spin glasses. In *Proc. Ninth International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC 2007)*, 415–418 (IEEE, 2007).
30. Oda, A. et al. Search for the ground states of ising spin clusters by using the genetic algorithms. *Int. J. Quantum Chem.* **105**, 645–654 (2005).
31. Ting, T., Yang, X.-S., Cheng, S. & Huang, K. Hybrid metaheuristic algorithms: past, present, and future. *Recent Advances in Swarm Intelligence and Evolutionary Computation* 71–83 (Springer Nature, 2015).
32. Bessarab, P. F., Uzdin, V. M. & Jónsson, H. Method for finding mechanism and activation energy of magnetic transitions, applied to skyrmion and antivortex annihilation. *Comput. Phys. Commun.* **196**, 335–347 (2015).
33. Furuya, A. et al. Semi-implicit steepest descent method for energy minimization and its application to micromagnetic simulation of permanent magnets. *IEEE Trans. Magn.* **51**, 1–4 (2015).
34. Ahn, C. W. & Ramakrishna, R. S. Elitism-based compact genetic algorithms. *IEEE Trans. Evolut. Comput.* **7**, 367–385 (2003).
35. Katoch, S., Chauhan, S. S. & Kumar, V. A review on genetic algorithm: past, present, and future. *Multimed. Tools Appl.* **80**, 8091–8126 (2021).

Acknowledgements

The authors acknowledge financial support from the Knut and Alice Wallenberg Foundation through Grant No. 2018.0060. Q.X. acknowledges the China Scholarship Council (201906920083), O.E. acknowledges eSENCE, the Swedish Research Council (VR) and the ERC (project FASTCORR–Grant No. 854843). A.D. acknowledges financial support from the Swedish Research Council (VR) through Grants No. 2019-05304 and 2016-05980. The computations/data handling were enabled by resources provided by the Swedish National Infrastructure for Computing (SNIC), partially funded by the Swedish Research Council through grant agreement No. 2016-07213. The authors would also like to extend our gratitude to Fredrik Dehlin, Liang Chen, Ivan Miranda, Filipp N. Rybakov, Heike Herper, and Liuzhen Yang for their valuable comments and discussions while writing our paper.

Author contributions

Q.X., O.E., and A.D. conceived the idea. Q.X., Z.S., P.H., O.E., and A.D. designed and performed the research. Q.X., Z.S., and P.H. developed the algorithm. M.P. and E.S.

discussed the results and commented on the manuscript. All authors contributed to discussions, writing, and revision of the manuscript.

Funding

Open access funding provided by Royal Institute of Technology.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-023-01360-4>.

Correspondence and requests for materials should be addressed to Qichen Xu or Anna Delin.

Peer review information *Communications Physics* thanks the anonymous reviewers for their contribution to the peer review of this work. A peer review file is available.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023