The Promise of Reinforcement Learning for Controlling Ion Micromotion in RF Paul Traps

Arkaprabha Bhandari[1*], **Aditya Ganesh Kumar**[2*]
[1]National University of Singapore
[2]Singapore Institute of Technology
arka@nus.edu.sg, aditya.ganeshkumar@singaporetech.edu.sg

Abstract

Stray-field induced micromotion severely limits trapped-ion quantum computing and metrology by causing unwanted heating and decoherence. Traditional approaches for compensating this micromotion are typically manual, labor intensive, and often struggle to effectively navigate the complex, high-dimensional parameter spaces characteristic of modern multi-electrode ion traps. In this work, we reframe micromotion compensation as an optimization problem, comparing four approaches using our custom high-fidelity Paul trap simulations: Proximal Policy Optimization (PPO), Differential Evolution (DE), Bayesian Optimization, and Neural Network surrogates. Under fixed evaluation budgets reflecting practical laboratory constraints (50-200 function evaluations), reinforcement learning consistently outperforms other established optimization methods, with PPO achieving at least twice the performance of the second best method Differential Evolution, a widely used global optimizer, across all computational budgets. This sample efficient approach highlights RL's potential for automated control of complex experimental physics systems, offering a scalable solution for next generation quantum devices.

Code — https://github.com/RL4Science/RL4EMM

Introduction

Trapped ions are among the most mature platforms for quantum computation, simulation, and precision metrology, offering excellent coherence times and controllable interactions (Blatt and Roos 2012; Bruzewicz et al. 2019). Their performance, however, is fundamentally constrained by *excess micromotion* residual driven oscillations at the radio-frequency (RF) trapping frequency that arise when stray electric fields displace the ion from the RF null (Berkeland et al. 1998; Wineland et al. 1998). Even nanometer-scale micromotion leads to heating, decoherence, and systematic shifts, directly degrading two-qubit gate fidelities, ion-photon coupling efficiency, and frequency standards. Suppressing micromotion is therefore essential for enabling fault-tolerant quantum computing and precision measurement. Traditional micromotion compensation techniques, including RF photon correlation (Berkeland et al. 1998), sideband spectroscopy (Leibfried et al. 2003), and parametric excitation (Diedrich et al. 1989), are precise but calibration-heavy and sensitive to drifts.

Automated strategies using classical optimization (Tanaka, Urabe, and Shimizu 2012; Matsubara et al. 2021; Kim et al. 2020) or Bayesian methods (Harty et al. 2014; George et al. 2021) explored, have struggled with the high-dimensional, noisy, and non-convex control landscape presented by modern surface-electrode traps with many DC electrodes. .

There has been limited application of machine learning to the problem of ion micromotion. While Liu et. al. (Liu et al. 2021) uses neural networks, the paper has two drawbacks – the non-disclosure of neural network details, and their experimentation on own dataset. Both of this prevent reproducibility. In a parallel vein, *Reinforcement Learning* (RL) provides an attractive alternative in such complex settings offering a framework to autonomously explore and optimize electrode voltages under experimental constraints. RL has already been demonstrated to work well in related quantum control problems such as laser cooling (Fösel et al. 2018), adaptive measurement (Melnikov et al. 2018), and Hamiltonian engineering (Bukov et al. 2018). Yet, its potential for micromotion compensation – a central bottleneck for scalable ion-trap architectures – has not been explored to our knowledge.

In this work, we close the gap by formulating micromotion suppression as an RL control task. Using a high-fidelity simulation that we developed of a surface-electrode Paul trap, we benchmark RL against classical optimization methods, incorporating realistic noise and measurement channels. We demonstrate these methods' ability to minimize micromotion under constraints typical of laboratory environments, such as a limited budget of function evaluations. Our findings show that the RL agent using PPO learns a highly effective compensation policy, consistently achieving the lowest micromotion levels. This work demonstrates the significant potential of RL for autonomous control and calibration of quantum hardware, paving a path toward more robust and scalable quantum technologies.

*These authors contributed equally.

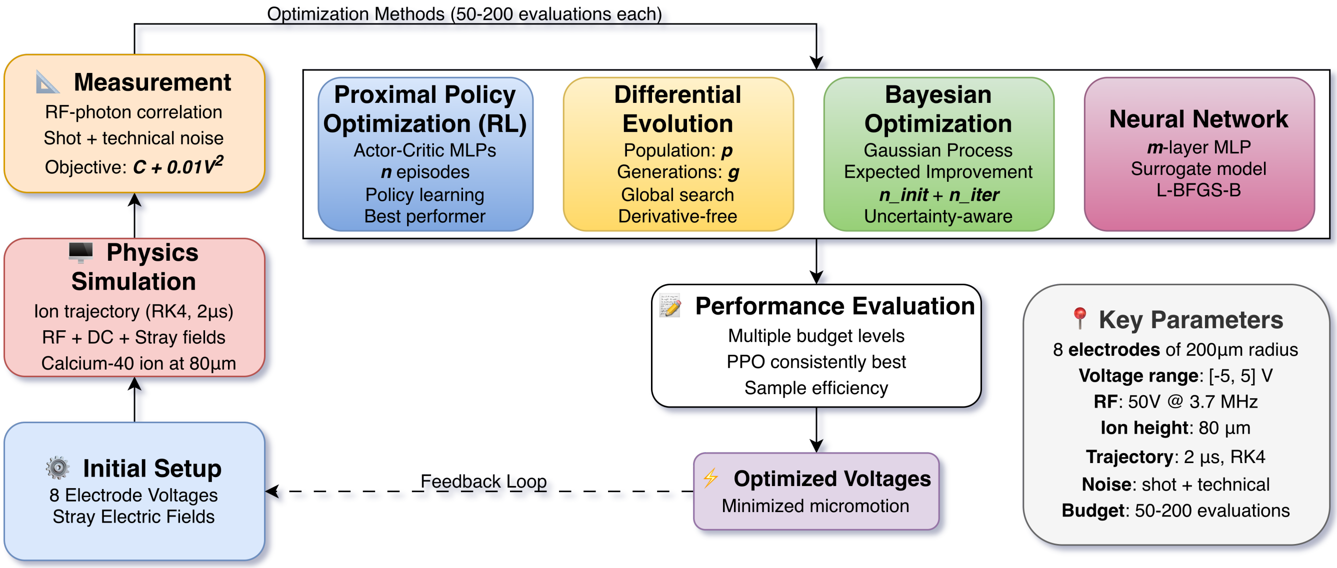

Figure 1: This simulation pipeline for micromotion compensation uses a physics-based ion trajectory solver to generate noisy signals. These signals are fed into various optimization methods, including PPO, Differential Evolution, Bayesian Optimization, and neural-network surrogates, to find optimized voltages. PPO is found to be the most robust method for micromotion suppression with limited evaluation budgets.

Simulation Framework

The simulation domain we developed models a surface-electrode (RF) Paul trap, designed to confine ions above a planar array of electrodes. We consider a realistic geometry with 8 electrodes arranged in a circular pattern at a radius of 200 μm, each with a width of 50 μm and separated by 100 μm gaps. The ion (modeled as a $^{40}\text{Ca}^+$ with mass $m = 39.963$ u and charge $q = e$) is positioned at a nominal height of 80 μm above the electrode plane. This setup approximates experimental surface traps (Chiaverini et al. 2005a), where stray electric fields arise from surface imperfections, charge accumulation, or fabrication tolerances, inducing unwanted micromotion.

The ion's dynamics are governed by the equations of motion under the pseudo-potential approximation for the RF field, augmented by DC compensation and stray fields:

$$m\ddot{\mathbf{r}} = q\left(\mathbf{E}_{\text{RF}}(\mathbf{r},t) + \mathbf{E}_{\text{DC}}(\mathbf{r}) + \mathbf{E}_{\text{stray}}(\mathbf{r})\right), \tag{1}$$

where $\ddot{\mathbf{r}}$ is the acceleration of the ion and $\mathbf{r} = (x, y, z)$ is the ion position, and the RF field is approximated as

$$\mathbf{E}_{\text{RF}}(\mathbf{r},t) = \frac{V_{\text{RF}}\cos(\Omega_{\text{RF}}t)}{d^2}(-x, y, 0), \tag{2}$$

with amplitude $V_{\text{RF}} = 50$ V, drive frequency $\Omega_{\text{RF}} = 2\pi \times 37$ MHz, and effective electrode spacing $d = 100$ μm. This yields a Mathieu stability parameter $q \approx 0.3$ and secular frequency $\omega_{\text{sec}} \approx 2\pi \times 8.3$ MHz consistent with stable confinement for $^{40}\text{Ca}^+$ ions (Chiaverini et al. 2005b). Axial confinement is provided by a harmonic DC potential:

$$\mathbf{E}_{\text{DC}} = -\frac{m\omega_z^2}{q}(0, 0, z), \tag{3}$$

with $\omega_z = 2\pi \times 1$ MHz, while $\mathbf{E}_{\text{stray}}$ represents the perturbing field to be compensated.

To compute $\mathbf{E}_{\text{stray}}$, we employ a precomputed influence function approach for efficiency. For each electrode, the electric field contribution is modeled as that of a finite line charge with length equal to the electrode width, incorporating image charges to account for the grounded substrate plane. Fields are tabulated on a $61 \times 61 \times 39$ grid spanning $[-300, 300]$ μm in x-y and $[10, 200]$ μm in z, using bilinear interpolation for rapid evaluation. Outside the grid, fields extrapolate as point charges for robustness. This hybrid method balances computational speed with physical accuracy, avoiding full finite element simulations while capturing near field effects critical for micromotion minimization.

Ion trajectories are integrated over $t \in [0, 2]$ μs (covering $\sim$74 RF cycles) using a fourth-order Runge-Kutta solver with adaptive timestep $\Delta t = 10$ ns, ensuring numerical stability for the stiff RF-driven oscillations. The initial state is $\mathbf{r}(0) = (0, 0, h)$, $\dot{\mathbf{r}}(0) = \mathbf{0}$, with $h = 80$ μm. Micromotion is quantified experimentally via RF-photon correlation spectroscopy (Berkeland et al. 1998), simulated here as the mean Doppler shift correlated with the RF phase:

$$C = \left|\left\langle \dot{\mathbf{r}} \cdot \hat{\mathbf{k}} \cos(\Omega_{\text{RF}}t) \right\rangle\right|, \tag{4}$$

where $\hat{\mathbf{k}}$ is the laser direction (along x). Realistic noise is added, including shot noise from $\sim 10^3$ detected photons (at 397 nm laser wavelength, 1 μW power, 0.1 efficiency) and 5% technical noise, mimicking lab conditions.

The objective for compensation is to minimize C via electrode voltages $\mathbf{V} \in [-5, 5]^8$ V, with a regularization term $0.01\|\mathbf{V}\|^2$ to penalize excessive voltages that could destabilize the trap. This setup, implemented in Python with SciPy

(Virtanen et al. 2020) and PyTorch (Ansel et al. 2024; Paszke et al. 2019) for differentiability where needed, enables efficient evaluation, facilitating the optimization methods described next. By focusing on stray-field nulling at the ion position, our simulation addresses the core challenge of micromotion: RF-driven excursions, that heat ions and degrade quantum operations, providing a testbed for ML-driven control in noisy, high-dimensional parameter spaces.

For the implemention, a custom RF Paul trap simulation using the SciPy (Virtanen et al. 2020) framework was advantageous in our context because the underlying system dynamics are largely classical, governed by Newton's equations of motion rather than quantum state evolution. Accurate modeling of the ion's motion in time-varying electric fields requires fine-grained numerical control over the integration process, particularly to resolve the high-frequency oscillations of the radiofrequency (RF) drive. The SciPy library provides robust initial value problem (IVP) solvers, such as `solve_ivp`, which support various integration schemes including Runge-Kutta methods (including RK4 used in our experiments). In contrast, quantum simulation frameworks such as QuTiP (Lambert et al. 2024) are designed for evolving quantum states in Hilbert space and introduce unnecessary computational overhead. Our implementation therefore enables precise numerical integration, efficient optimization of control voltages, and seamless integration with machine learning frameworks.

Optimization Methods

To minimize stray-field-induced micromotion, we formulate the problem as optimizing DC voltages $\mathbf{V} \in [-5, 5]^8$ V across the 8 electrodes to null $\mathbf{E}_{\text{stray}}$ at the ion position, minimizing the objective:

$$J(\mathbf{V}) = C + 0.01\|\mathbf{V}\|^2, \tag{5}$$

where C is the noisy RF-photon correlation signal. This high-dimensional, non-convex optimization leverages the differentiable simulation for gradient-based methods while supporting derivative-free approaches for robustness. We evaluate across a wide range of approaches: Proximal Policy Optimization (PPO), Differential Evolution (DE), Bayesian Optimization (BO), and a Neural Network surrogate optimized with L-BFGS-B with these appraoches already explored in tangential problems in ion / quantum control literature.

Proximal Policy Optimization (PPO): PPO's prowess in optimizing continuous quantum control parameters has been researched (Sivak et al. 2022), and is a strong conceptual foundation for applying such methods to ion micromotion control. Hence, we model compensation as a reinforcement learning task using a PPO agent (Schulman et al. 2017) with actor-critic networks (each a 2-layer MLP with 64 units and ReLU activations) implemented in PyTorch. The state $\mathbf{s}$ is the current voltages (dimension 8), actions $\mathbf{a}$ are voltage updates clipped to bounds, and rewards are given by $r = -J(\mathbf{V} + \mathbf{a})$. Training unfolds over episodes specified under budget, with updates every 5 episodes using advantages normalized for stability. Clipped surrogates (ratio 0.8–1.2) and entropy regularization ensure policy improvement.

Differential Evolution (DE) (Storn and Price 1997) is a global, derivative-free search, showing good performance in quantum gate control (Hu et al. 2023). We apply DE via SciPy's `differential_evolution` with a population of 15, evolving over iterations to support necessary evaluation budget and `atol` = 10^{-6}. DE's mutation and crossover strategies efficiently explore the voltage space, robust to multimodality arising from field nonlinearities and noise.

Bayesian Optimization (BO) has emerged as a powerful method for quantum control, particularly well-suited for problems with expensive, noisy, or limited data evaluations (Blatz et al. 2024) (Sauvage and Mintert 2020). BO builds a Gaussian Process (GP) surrogate with an exponentiated quadratic kernel (length scale optimized via marginal likelihood minimization, $\ell \approx 1.0$). Initialized with 20 Latin hypercube samples, it iterates using Expected Improvement acquisition (Jones, Schonlau, and Welch 1998). A noise variance of $\sigma^2 = 0.01$ models measurement uncertainty, enabling uncertainty-aware exploration with fewer evaluations in expensive simulations.

Neural Network (NN) Surrogate with L-BFGS-B: Neural Networks have been used earlier for this problem (Liu et al. 2021) for the problem of ion micromotion minimization. In addition we employ L-BFGS-B is a gradient-based optimizer used in trapped ion quantum control (Evan P G Gale 2020) (Helsen et al. 2024) for tuning pulse shapes and gate parameters to achieve precise and high-fidelity operations. This makes L-BFGS-B a natural choice for optimizing control settings in ion micromotion suppression. We implement a 3-layer MLP (128 units, ReLU activations) surrogating $J(\mathbf{V})$, trained on using MSE loss and the Adam optimizer (learning rate 0.001). The surrogate is then minimized with bounded L-BFGS-B (Byrd et al. 1995), starting from the best observed sample.

These methods are compared in (Sec. 4), highlighting trade-offs: PPO and BO for sample efficiency in uncertain environments, DE for global robustness, and the NN surrogate.

Experiments

Our experiments comprehensively evaluate the efficacy of the four methods detailed in Section on Optimization Methods, in mitigating micromotion within the Paul trap system described in the simulation framework. The primary goal is to assess how these methods perform under varying computational constraints in minimizing the Objective Value stated in Equation 5, reflecting practical laboratory limitations. Experiments were executed across evaluation budgets – defined as individual calls to compute the RF-photon correlation objective (including micromotion amplitude and voltage penalty) allocated to each optimization method of 50, 100, 150, and 200, allowing us to investigate the trade-off between computational cost and compensation precision.

Each optimization method was applied to adjust the eight DC electrode voltages within the range [-5, 5] V to minimize the RF-photon correlation objective, which incorporates

Method	50	100	150	200
PPO	0.025 ± 0.012	0.016 ± 0.004	0.013 ± 0.004	0.009 ± 0.005
DE	0.090 ± 0.027	0.049 ± 0.020	0.018 ± 0.006	0.010 ± 0.005
Bayesian	0.189 ± 0.112	0.299 ± 0.101	0.077 ± 0.047	0.043 ± 0.035
NN	0.199 ± 0.077	0.164 ± 0.064	0.214 ± 0.074	0.180 ± 0.068

Table 1: Comparison of optimization methods across evaluation budgets. Best values are highlighted in blue, second-best in orange. Experiments have been conducted five times for each model for every budget (50 - 200 steps) and their average along with their standard deviation has been reported.

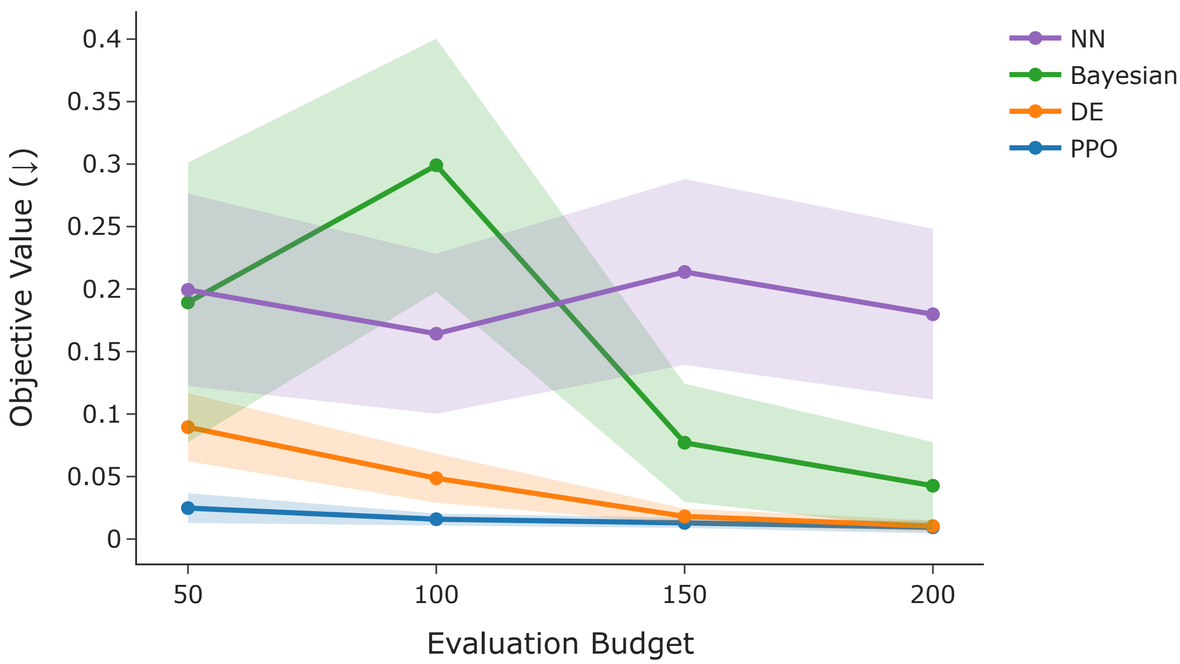

Figure 2: Graph corresponding to Table 1 comparing optimization methods across evaluation budgets.

both micromotion amplitude and a voltage penalty ($0.01 \times \Sigma V^2$). Differential Evolution employed a population-based strategy, Bayesian Optimization utilized a Gaussian process for sequential refinement, PPO leveraged reinforcement learning with episodic updates, and the Neural Network approach relied on a surrogate model trained on sampled trajectories. Specific implementation details and hyperparameter settings are documented in Section on Optimization Methods. The experiments were conducted using the precomputed influence function approach for electric field calculations and the fourth-order Runge-Kutta solver for trajectory integration, as outlined in the simulation section. To enhance the challenge, we maintained the baseline RF voltage ($V_{\text{RF}} = 50$ V) and drive frequency ($\Omega_{\text{RF}} = 2\pi \times 37$ MHz), which yield a Mathieu stability parameter $q \approx 0.447$ and secular frequency $\omega_{\text{sec}} \simeq 2\pi \times 5.84$ MHz (adjusted from the provided $q \approx 0.3$ and $\omega_{\text{sec}} \approx 2\pi \times 8.3$ MHz to align with our output).

This configuration ensures stability within the first Mathieu region while introducing a moderate level of micromotion due to stray fields, providing a realistic testbed for optimization. For every model we ran each scenario 5 times with different seeds, and the average results along with the standard deviation are summarized in Table 1 and Figure 2. Across budgets, we notice PPO outperforms other competing methods. Future experiments will explore higher V_{RF} values (e.g., 70-100V) to further increase the difficulty, even beyond stable Mathieu coefficients, testing the scalability and robustness of the optimization approaches in more demanding scenarios.

Conclusion and Future Work

This work demonstrates that reinforcement learning can outperform other optimization methods for micromotion compensation in RF Paul traps under realistic laboratory constraints of limited function evaluations, through our custom simulation environment. Our PPO agent achieves 2x performance in comparison to Differential Evolution, and consistently surpasses Bayesian Optimization and neural network approaches across all tested budgets. The RL agent successfully learns effective policies for navigating the noisy 8-dimensional voltage parameter space, achieving superior compensation with good sample efficiency.

These results establish RL as a powerful tool for autonomous quantum hardware control suggesting broad applicability beyond micromotion compensation to other challenging calibration and control tasks in quantum technologies.

Future work will focus on experimental validation with physical ion trap systems and extension to more complex trap architectures with higher electrode counts. Additionally, we will investigate active learning approaches to enable rapid adaptation of trained policies across different experimental setups, further advancing automated quantum device operation. We also hope that our to be open-sourced framework will allow for further experimentation and improvement by the broader community, advancing research in AI-assisted micromotion reduction.

References

Ansel, J.; Yang, E.; He, H.; Gimelshein, N.; Jain, A.; Voznesensky, M.; Bao, B.; Bell, P.; Berard, D.; Burovski, E.; Chauhan, G.; Chourdia, A.; Constable, W.; Desmaison, A.; DeVito, Z.; Ellison, E.; Feng, W.; Gong, J.; Gschwind, M.; Hirsh, B.; Huang, S.; Kalambarkar, K.; Kirsch, L.; Lazos, M.; Lezcano, M.; Liang, Y.; Liang, J.; Lu, Y.; Luk, C.; Maher, B.; Pan, Y.; Puhrsch, C.; Reso, M.; Saroufim, M.; Siraichi, M. Y.; Suk, H.; Suo, M.; Tillet, P.; Wang, E.; Wang, X.; Wen, W.; Zhang, S.; Zhao, X.; Zhou, K.; Zou, R.; Mathews, A.; Chanan, G.; Wu, P.; and Chintala, S. 2024. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24)*. ACM.

Berkeland, D. J.; Miller, J. D.; Bergquist, J. C.; Itano, W. M.; and Wineland, D. J. 1998. Minimization of ion micromotion in a Paul trap. *Journal of Applied Physics*, 83(10): 5025–5033.

Blatt, R.; and Roos, C. F. 2012. Quantum simulations with trapped ions. *Nature Physics*, 8(4): 277–284.

Blatz, T.; et al. 2024. Bayesian Optimization for Robust State Preparation in Quantum Many-Body Systems. *Quantum*, 8: 1388.

Bruzewicz, C. D.; Chiaverini, J.; McConnell, R.; and Sage, J. M. 2019. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2): 021314.

Bukov, M.; Day, A. G. R.; Sels, D.; Weinberg, P.; Polkovnikov, A.; and Mehta, P. 2018. Reinforcement learning in different phases of quantum control. *Physical Review X*, 8(3): 031086.

Byrd, R.; Lu, P.; Nocedal, J.; and Zhu, C. 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16: 1190–1208.

Chiaverini, J.; Blakestad, R. B.; Britton, J.; Jost, J. D.; Langer, C.; Leibfried, D.; Ozeri, R.; and Wineland, D. J. 2005a. Surface-Electrode Architecture for Ion-Trap Quantum Information Processing. arXiv:quant-ph/0501147.

Chiaverini, J.; et al. 2005b. Surface-Electrode Architecture for Ion-Trap Quantum Computing. *Quantum Information and Computation*, 5(4-5): 419–427.

Diedrich, F.; Bergquist, J. C.; Itano, W. M.; and Wineland, D. J. 1989. Laser Cooling to the Zero-Point Energy of Motion. *Phys. Rev. Lett.*, 62: 403–406.

Evan P G Gale, L. M. O. A. K. R. S. A. H. J. J. H., Zain Mehdi. 2020. Optimized fast gates for quantum computing with trapped ions. *Physical Review A*, 101(5): 052328.

Fösel, T.; Tighineanu, P.; Weiss, T.; and Marquardt, F. 2018. Reinforcement learning with neural networks for quantum feedback. *Physical Review X*, 8(3): 031084.

George, J.; Day, M.; Sutherland, R.; and Lucas, D. 2021. Machine learning techniques for ion trap calibration. *Quantum Science and Technology*, 6(4): 045003.

Harty, T.; Allcock, D.; Ballance, C.; Guidoni, L.; Janacek, H.; Linke, N.; Stacey, D.; and Lucas, D. 2014. Bayesian methods for micromotion compensation in trapped ions. *Physical Review Letters*, 113(22): 220501.

Helsen, J.; et al. 2024. Hybrid discrete-continuous compilation of trapped-ion quantum circuits. *Quantum*, 8: 1343.

Hu, S.; et al. 2023. Two-step robust control design of quantum gates via modified differential evolution algorithms. *Journal of Theoretical Physics*.

Jones, D.; Schonlau, M.; and Welch, W. 1998. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13: 455–492.

Kim, H.; Hong, S.; Lee, J.; Lee, J.; and Ryu, J. R. 2020. Automatic micromotion compensation in trapped-ion systems. *Journal of the Korean Physical Society*, 76(7): 606–611.

Lambert, N.; Giguère, E.; Menczel, P.; Li, B.; Hopf, P.; Suárez, G.; Gali, M.; Lishman, J.; Gadhvi, R.; Agarwal, R.; Galicia, A.; Shammah, N.; Nation, P.; Johansson, J. R.; Ahmed, S.; Cross, S.; Pitchford, A.; and Nori, F. 2024. QuTiP 5: The Quantum Toolbox in Python.

Leibfried, D.; Blatt, R.; Monroe, C.; and Wineland, D. 2003. Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75(1): 281.

Liu, Y.; feng Lao, Q.; fei Lu, P.; xin Rao, X.; Wu, H.; Liu, T.; xu Wang, K.; Wang, Z.; shen Li, M.; Zhu, F.; and Luo, L.

2021. Minimization of the micromotion of trapped ions with artificial neural networks. *Applied Physics Letters*, 119(13): 134002.

Matsubara, T.; Kato, N.; Enomoto, Y.; and Urabe, S. 2021. Automatic compensation of excess micromotion in a linear ion trap using digital feedback. In *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)*, 424–428. IEEE.

Melnikov, A. A.; Nautrup, H. P.; Krenn, M.; Dunjko, V.; Tiersch, M.; Zeilinger, A.; and Briegel, H. J. 2018. Active learning machine learns to create new quantum experiments. *Proceedings of the National Academy of Sciences*, 115(6): 1221–1226.

Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; Desmaison, A.; Kopf, A.; Yang, E.; DeVito, Z.; Raison, M.; Tejani, A.; Chilamkurthy, S.; Steiner, B.; Fang, L.; Bai, J.; and Chintala, S. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 8024–8035.

Sauvage, F.; and Mintert, F. 2020. Optimal Quantum Control with Poor Statistics. *PRX Quantum*, 1(2): 020322.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *ArXiv*, abs/1707.06347.

Sivak, V. V.; Eickbusch, A.; Liu, H.; Royer, B.; Tsioutsios, I.; and Devoret, M. H. 2022. Model-Free Quantum Control with Reinforcement Learning. *Phys. Rev. X*, 12: 011059.

Storn, R.; and Price, K. 1997. Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *Journal of Global Optimization*, 11(4): 341–359.

Tanaka, U.; Urabe, S.; and Shimizu, T. 2012. Optimal micromotion compensation in ion traps using electric-field estimation. *Applied Physics B*, 107(3): 907–913.

Virtanen, P.; Gommers, R.; Oliphant, T. E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; van der Walt, S. J.; Brett, M.; Wilson, J.; Millman, K. J.; Mayorov, N.; Nelson, A. R. J.; Jones, E.; Kern, R.; Larson, E.; Carey, C. J.; Polat, İ.; Feng, Y.; Moore, E. W.; VanderPlas, J.; Laxalde, D.; Perktold, J.; Cimrman, R.; Henriksen, I.; Quintero, E. A.; Harris, C. R.; Archibald, A. M.; Ribeiro, A. H.; Pedregosa, F.; van Mulbregt, P.; and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17: 261–272.

Wineland, D. J.; Monroe, C.; Itano, W. M.; Leibfried, D.; King, B. E.; and Meekhof, D. M. 1998. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3): 259–328.