



OPEN

Dynamic compensation of stray electric fields in an ion trap using machine learning and adaptive algorithm

Moji Ghadimi¹✉, Alexander Zappacosta¹, Jordan Scarabel¹, Kenji Shimizu¹, Erik W. Streed^{1,2} & Mirko Lobino^{1,3}

Surface ion traps are among the most promising technologies for scaling up quantum computing machines, but their complicated multi-electrode geometry can make some tasks, including compensation for stray electric fields, challenging both at the level of modeling and of practical implementation. Here we demonstrate the compensation of stray electric fields using a gradient descent algorithm and a machine learning technique, which trained a deep learning network. We show automated dynamical compensation tested against induced electric charging from UV laser light hitting the chip trap surface. The results show improvement in compensation using gradient descent and the machine learner over manual compensation. This improvement is inferred from an increase of the fluorescence rate of 78% and 96% respectively, for a trapped $^{171}\text{Yb}^+$ ion driven by a laser tuned to -7.8 MHz of the $^2\text{S}_{1/2} \leftrightarrow ^2\text{P}_{1/2}$ Doppler cooling transition at 369.5 nm.

Trapped ions are one of the most promising candidates for implementing scalable quantum information processing^{1–4}. Traditionally macroscopic rod traps were used for electrically trapping laser cooled ions. These Paul traps have a deep potential well that enables the trapping of multiple ions in a chain, but their distant electrodes lack the fine spatial control of electric fields necessary to efficiently split and combine small ion crystals⁵. One of the proposed architectures for scaling up the number of trapped ions in a quantum processor is the use of microfabricated surface chip traps with multiple DC electrodes that are able to manipulate multiple ions individually^{6–10}.

In these chip traps the ions are confined to a potential formed by the node of an oscillating RF electric field in two dimensions while a DC field generated from multiple electrodes provides a finely spatially adjustable potential in the third dimension. It is well-known in Paul trapping that stray DC electric fields push the ion off the RF node and induce micromotion that degrades or outright prevents effective confinement and laser-cooling¹¹. The solution to this problem is to use the DC electrodes to generate an opposite DC electric field that compensates for these stray fields⁹.

Micromotion amplitude can be measured and dealt with in multiple ways^{9,12–15}. Most of the conventional methods are time consuming and can only be done manually. A simple and quick proxy, that is only applicable to small amounts of micromotion, is to measure the ion's fluorescence under laser cooling near the Doppler limit. In fact for low magnitudes of micromotion, if a laser is red-detuned near the natural linewidth of the atom, the ion fluorescence rate increases as the micromotion amplitude decreases (see “[Experimental setup](#)” section and Fig. 2a). While stray electric fields can be readily compensated in simple fixed trap geometries¹⁶, this task becomes more challenging for multi-electrode designs¹⁷ and in proximity of surfaces which are vulnerable to laser induced charging¹⁸.

Automation, optimization, and machine learning have been used to improve different manual tasks in atom and ion traps^{19–21} and they are also useful tools for optimizing the individual electrode voltages which generate the trapping electric field. In this work we show how the voltages of an array of electrodes in a surface chip trap can be optimized using a fully automated, machine learning driven process to minimize micromotion and maximize fluorescence. The chip trap in our experiment has 44 segmented DC electrodes and 2 RF rails, as shown in Fig. 1. The trap incorporates nearly aberration free diffractive mirrors directly fabricated onto the

¹Center for Quantum Dynamics, Griffith University, Nathan, QLD, Australia. ²Institute for Glycomics, Griffith University, Southport, QLD, Australia. ³Queensland Micro Nanotechnology Centre, Nathan, QLD, Australia. ✉email: moji131@gmail.com

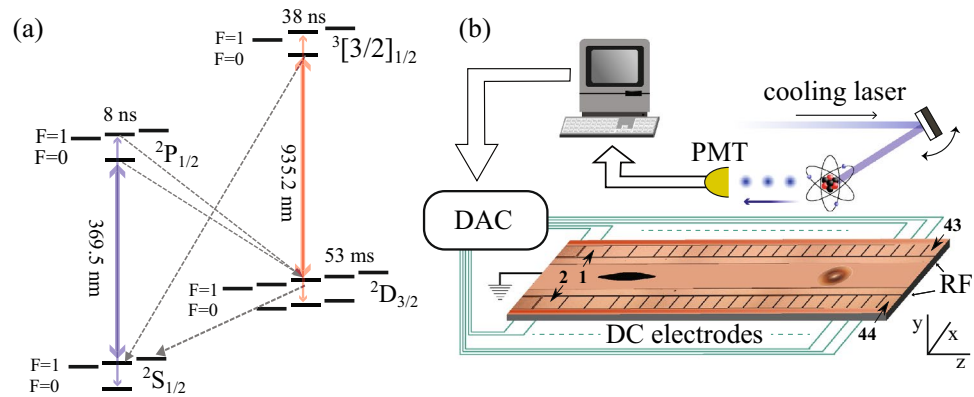


Figure 1. (a) Transitions for laser cooling $^{171}\text{Yb}^+$. The ion is laser cooled on the nearly closed 369.5 nm $^2\text{S}_{1/2}$, $F = 1$ to $^2\text{P}_{1/2}$, $F = 0$ transition. The small amount of off-resonant scattering into the $^2\text{S}_{1/2}$, $F = 0$ state is repumped with a 14.7 GHz sideband on the primary cooling laser. Decays from $^2\text{P}_{1/2}$ can also populate the $^2\text{D}_{3/2}$ level, which are repumped using a 935 nm laser. (b) Schematics of the ion trap chip and the optimization devices. RF and DC electrodes are used to trap the ion and shuttle it above the diffracting mirror that collects photons from the ion and sends it to photo multiplier tube (PMT). The PMT counts are processed by the optimiser and DC electrode voltages and cooling laser position are updated using a digital to analogue converter (DAC).

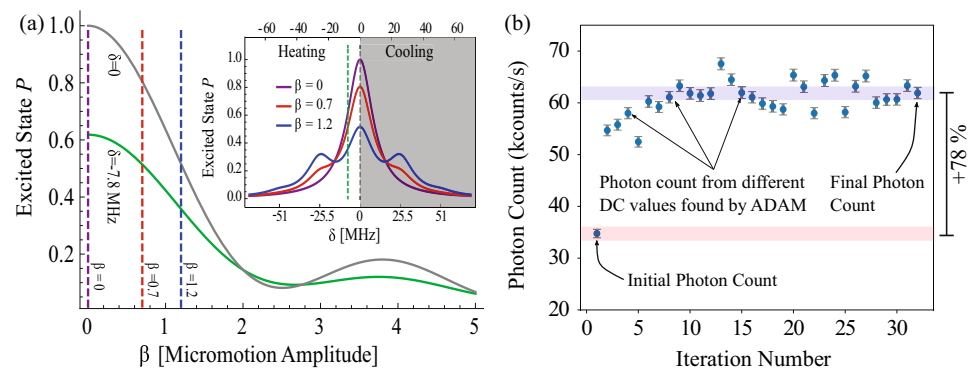


Figure 2. (a) Ideal micromotion broadening of the transition fluorescence in a two-level system as a function of the micromotion amplitude [β] and laser detuning (δ) in the low intensity limit (Eq. 1). For our 370 nm transition the fluorescence rate is proportional to the $^2\text{P}_{1/2}$ excited state population. The laser frequency detuning $\delta = -7.8$ MHz, the natural linewidth $\gamma = 20$ MHz and the RF trap drive frequency $\Omega = 25.5$ MHz. Heating/cooling dynamics are neglected. Our optimizations occur in the low beta (< 1) limit. (b) The photon count improvement using ADAM starting from manually optimized voltage set weights. This graph demonstrates a $78 \pm 1\%$ improvement in ion fluorescence. The error bars are a combination of the readout error of the PMT and peak to peak variations in the photon count whilst not optimizing. The background photon count here is 1184 ± 34 counts/s.

central ground electrode for efficient fluorescence collection from the ion²². A Gradient descent (ADAM²³) and a machine learning algorithm (MLOOP²⁰) were compared on the basis of the versatility and time taken to find the optimal compensation and highest fluorescence rate of the ion. These methods were applied to a trapped $^{171}\text{Yb}^+$ ion and compared to a manual optimization performed by an experienced operator.

With ADAM we were able to improve the fluorescence by 78% starting from voltages already optimized by manual adjustment, while MLOOP achieved a 96% improvement. Subsequently we tested the versatility and adaption of this procedure by deliberately charging the trap with UV light to drop the fluorescence rate by around 35% and compensating back to the optimal with ADAM. This was not tested on MLOOP as it was observed to be quite sensitive to the noise in the fluorescence signal and did not make a solid candidate for final optimization of the artificially charged trap.

Experimental setup

The ion trap used in these experiments and described in^{22,24}, consists of a planar rectangular chip of $1400 \mu\text{m}$ length and $800 \mu\text{m}$ width. It has a set of aluminum electrodes patterned on top and is schematically shown in Fig. 1. The ion is trapped $60 \mu\text{m}$ above the surface, between the two RF electrode rails that extend for the full length of chip, and 44 DC electrodes on the sides. The DC electrodes are used for creating a reconfigurable

trapping potential along the length of the trap, compensating stray electric fields, and enabling ion crystals merging, separation, and shuttling.

An atomic oven underneath one end of the trap generates a beam of neutral Yb atoms which passes through a slit in the chip. The ^{171}Yb is first excited by an isotopically selective 399 nm laser²⁵, and subsequently non-resonantly ionized by a 369.5 nm laser. The ion is then Doppler cooled using the same 369.5 nm laser tuned nearly resonant with $^2\text{S}_{1/2} F = 1 \leftrightarrow ^2\text{P}_{1/2} F = 0$ transition. To efficiently cool the ion in both x and y directions, the RF axis of the trap is slightly tilted and therefore the horizontal cooling laser has a component along both nearly horizontal x-axis and nearly vertical y-axis. Occasionally off-resonant scattering from the $^2\text{P}_{1/2} F = 1$ state will populate the dark $^2\text{S}_{1/2} F = 0$ ground state, which is repumped by a small amount of light from a 14.7 GHz sideband added to the primary 369.5 nm cooling laser. There is also a probability of 0.5% for the atom to decay from the $^2\text{P}_{1/2}$ into the meta-stable $^2\text{D}_{3/2}$ state. In this case a 935 nm laser repumps the ions back into the cooling cycle. Figure 1a depicts the cooling and repumping transitions.

After initial trapping near the oven slit region, the ion is shuttled along the length of the trap by properly controlling the voltages of the array of DC electrodes with a 12-bit National Instruments PXI-6713 DAQ (output doubled to $\pm 20\text{V}$ with 10 mV resolution). For our experiments, the ion was shuttled above the surface of an integrated diffracting micro-mirror patterned on the ground electrode with a focal length of $60\ \mu\text{m}$, equal to the height of the ion above the trap, to improve fluorescence collection and coupling into a single mode fiber²².

For efficient Doppler cooling and subsequent quantum information processing experiments, it is important to minimize any micromotion induced by stray electric fields. This is done by tuning the voltages of the DC electrodes so that they generate an opposing electric field that compensates for the stray fields. To manually compensate stray electric fields in the RF confinement plane xy (see Fig. 1), it is not practical to individually tune all the 44 DC electrodes. We instead use two voltage sets (array of DC voltages for individual electrodes) which were calculated for generating an electric field of 100 V/m in the x (horizontal) and y (vertical) directions. Electrode voltages of -1 to $+1\text{V}$ for the x direction and -2 to $+2\text{V}$ for the y direction are needed to create electric field of 100 V/m at the location of the ion. These sets of voltages are multiplied by arbitrary weights w_x and w_y to compensate arbitrary fields in the RF trapping plane. In practice this process has imperfections and cannot fully compensate for the stray electric fields¹⁷. Instead, individual electrodes need to be tuned independently for optimal compensation. This optimum compensation is position dependent and there might be multiple solutions for a multi-electrode trap.

Two other DC voltage sets are used and their relative weights tuned: a harmonic voltage set, and RF plane trap axis rotation voltage set. The former creates a harmonic potential to trap the ion in different locations along the chip length, while the latter rotates the trap's secular confinement axes to lift the degeneracy and improve cooling²⁶.

We infer the magnitude of the stray electric field from the fluorescence rate of the ion when the cooling laser is red-detuned near half its natural linewidth for optimal cooling. When driving an ideal cold two-level atom with this laser, the fluorescence is proportional to the population of excited state, P_e , that can be calculated from the optical Bloch equation²⁷. The relationship between P_e and the strength of the stray electric field is given by¹⁵:

$$P_e(\beta) = C \sum_{m=-\infty}^{\infty} \frac{J_m^2(\beta)}{\left(\frac{\delta_L}{\gamma} + m\frac{\Omega}{\gamma}\right)^2 + \frac{1}{4}} \quad (1)$$

where J_m is the m th order Bessel function of the first kind, $\delta_L = \omega_{atom} - \omega_{laser}$ is the laser frequency detuning, γ is the natural linewidth and Ω is the RF trap frequency. The amount of micromotion is embedded in the dimensionless parameter β which accounts for the dimensionless measure of the ion's coherently driven motion from the stray electric fields and RF phase imbalances (AC micromotion)¹⁵. C depends on the strength of the coupling between the levels that is a constant here since we keep laser intensity and direction fixed. Figure 2a shows a graph of $P_e(\beta)/P_e(0)$ as a function of β . The graph shows that at our detuning (-7.8 MHz) the population decreases monotonically as β increases if β is less than 2.5. For larger β micromotion induced local maxima arise that prevent inferring the magnitude of micromotion from fluorescence rate. Since our operating regime is below this threshold, we can use the change in fluorescence rate to detect the change in magnitude of the stray electric field.

Results

A gradient descent algorithm (ADAM) and a deep learning network (MLOOP) were tested for compensating stray fields in different working regimes. The source code used for the experiments is available in²⁸. The software controlled the voltages using the PXI-6713 DAQ and read the fluorescence counts from a photo-multiplier-tube (PMT) through a time tagging counter (IDQ id800). All software was written in python and interfaced with the DAQ hardware using the library NI-DAQmx Python. A total of 44 DC electrodes and the horizontal position of the cooling laser were tuned by the program, resulting in a total of 45 input parameters.

Gradient descent optimizer. The first compensation test was performed by ADAM gradient descent algorithm. This is a first order optimizer that uses the biased first and second order moments of the gradient to update the inputs of an objective function, and was chosen for its fast convergence, versatility in multiple dimensions and tolerance to noise²³. Our goal was to maximize the fluorescence of the ion which was described by a function $f(\vec{\alpha})$, where $\vec{\alpha} = (\alpha_1, \alpha_2, \alpha_3, \dots, \alpha_{45})$ represents the array of parameters to be optimized. To find the optimal $\vec{\alpha}$, the algorithm needs to know the values of the partial derivatives for all input parameters. Because we do not have an analytic expression for $f(\vec{\alpha})$, the values of its derivatives were estimated from experimental measurements by sequentially changing each input α_i , and reading the associated change in fluorescence f . This data were used as inputs to ADAM for finding the optimal $\vec{\alpha}$ which maximized f .

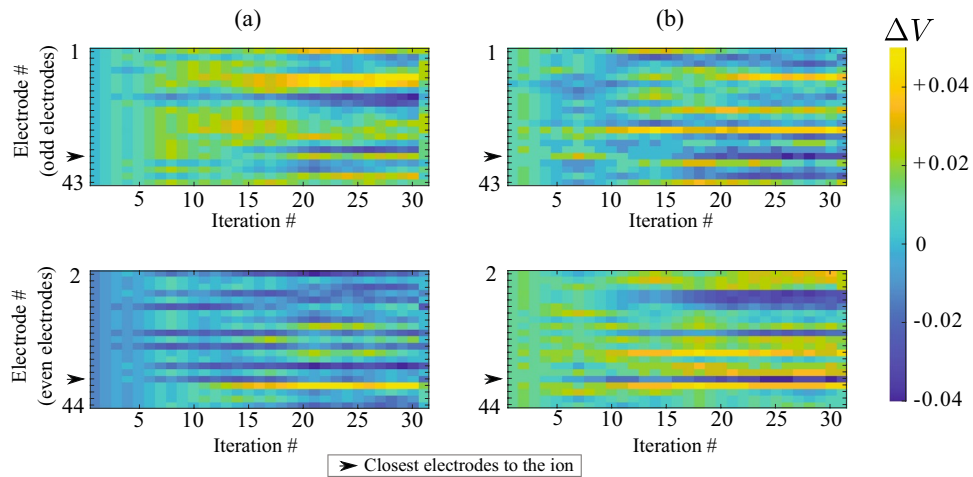


Figure 3. Voltage deviation from the original starting point during optimization with ADAM. (a) uncharged trap (“Gradient descent optimizer” section). (b) During UV charging (“Testing under poor trap conditions”). Top graphs show odd electrode numbers corresponding to top DC electrodes in Fig. 1b and the bottom graphs show the even electrode numbers. The values were determined by subtracting the voltage at each iteration by the starting voltage $\Delta V = V_n - V_0$. Changes can be seen in almost all the electrodes of the trap.

Before running the automated compensation, we manually adjusted the 4 weights of the voltage sets used for compensation described in the previous section. We also tried to run ADAM to optimize these 4 parameters but the increase in fluorescence was limited to 6%. After manual compensation, we ran ADAM on all 45 inputs with the algorithm parameters given in the source code²⁸. Each iteration took 12 s, where 9.8 s were the photon readout (0.1 s \times 2 readouts per parameter plus 2 \times 0.1 s readouts at the beginning and end of the iteration), and the rest of the time was the gradient computation. If the photon count reduced by more than 40% of its initial value, the algorithm terminated and applied the previously found optimum. This acted as the safety net for the program, ensuring the ion was not lost while optimizing the 45 inputs. We need this safety net because if the ion is heated past the capture range for the used cooling detuning, it will be ejected from the trap. In our implementation of the algorithm we removed the reduction in the step size of the optimization algorithm as iterations progressed. This step reduction, which is present in the standard version of ADAM, is not ideal when stray fields change with time since the optimal values of the voltages also drift in time. The removal caused some fluctuations in the photon readout near the optimal settings. Adding to these fluctuations, other sources of noise, such as wavemeter laser locking³², and mechanical drift in the trap environment, resulted in daily photon count variations of around 5%. Fluctuation in laser power was not a concern here since the power of the cooling laser was stabilized. Despite these fluctuations, and the fact that stray fields change every day, the algorithm demonstrated an increase in fluorescence collection up to $78 \pm 1\%$ (Fig. 2b) when starting from a manually optimized configuration in less than 10 iterations, or 120 s.

The ADAM algorithm was fast and reliable (the ion was never lost during optimization), even in extremely volatile conditions like having time-dependent charging and stray electric field buildup. Figure 3a shows a colourmap of the voltages and laser position adjustments, where most of the improvement came from adding the same voltage to all DC electrodes indicating that the ion was not at optimal height. The volatility of the ion-trap environment causes the fluorescence rate to oscillate around the optimal point. To get the best value, instead of using the values of the final iteration, the software saved all voltage combinations and applied the setting with the highest photon count after all iterations were finished. Despite picking the best value it can be seen in Fig. 2b that the fluorescence for some iterations during the optimization are higher than the final point selected by the software. This is because when the settings are changed, the ion fluorescence rate may transiently increase and subsequently stabilize to a slightly lower value for the same voltage settings.

Deep learning network for the ion trap. The second algorithm tested was a deep learning network using the python based optimization and experimental control package MLOOP²⁰. MLOOP uses Differential Evolution³³ for exploring and sampling data. The blue points in Fig. 4a corresponds to these samples and it can be seen that even at the end of optimization, they can have non-optimum fluorescence rates. MLOOP also trains a neural network using the data collected by Differential Evolution and creates an approximate model of the experimental system. It then uses this model to predict an optimum point. The red points in Fig. 4a shows the optimum points predicted by the neural network model. It can be seen that this section starts later than Differential Evolution, as it requires some data for initial neural network training, and gradually finds the optimum and stays near it. For training of the neural network, the inbuilt ADAM optimizer is used to minimize the cost function. The sampling in MLOOP does not require a gradient calculation which greatly improves the sampling time. Even though the sampling is fast, training the network to find an optimal point requires a minimum of 100 samples and that makes MLOOP slower than ADAM. With our settings for MLOOP, each iteration took 0.7 s on average and therefore 700 s was needed to take 1000 samples shown in Fig. 4a.

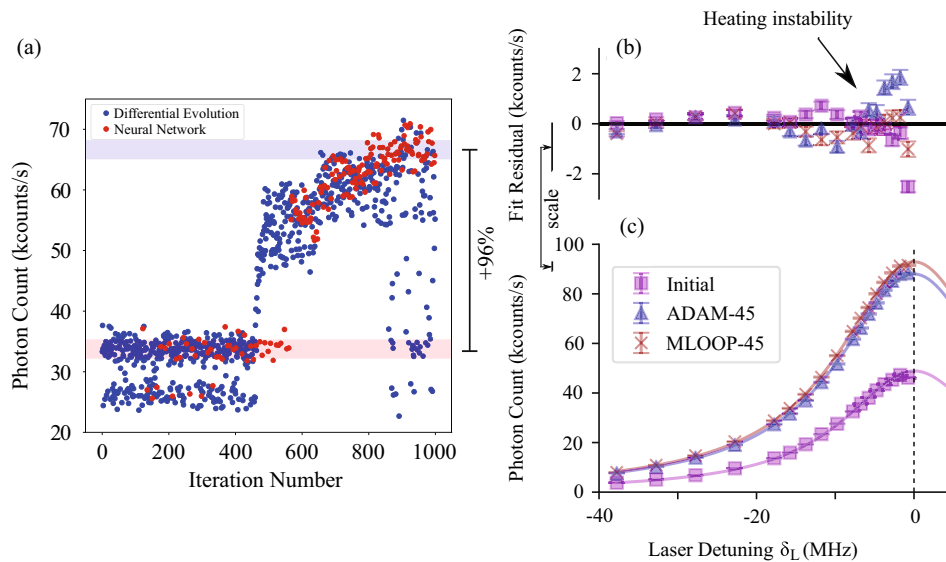


Figure 4. (a) MLOOP deep learning network. Differential Evolution explores the input space (blue points) and the neural network creates a model of the data and predicts an optimum (red points). Maximum photon count of the neural network points is $96 \pm 1\%$ higher than manual optimization. Differential Evolution continues to explore the input space and has varied photon counts. The beginning point for the process (found by manually adjusting the 4 voltage set weights) was at 33700 counts/s and the highest photon count found by the neural network was at 66200 counts/s. (b) and (c) Fluorescence versus laser frequency detuning from the resonance for initial setting and after different optimizations. It can be seen from that the experimental values are very close to the theoretical Lorentzian fit^{29–31}. This shows the heating is low before and after optimization and therefore the change in fluorescence can be used to infer the change in heating. Deviation from the theory near the resonance shown in (b) is a sign of small heating instability.

In our test the neural network in MLOOP had 5 layers with 45 nodes each, all with Gaussian error correction. The neural network structure (number of layers and cells) was manually optimized and tested on a 45-dimensional positive definite quadratic function before being used for the experiment. Once the ion was trapped, positioned above the integrated mirror²², and photon counts were read, the program started sampling 100 different voltage combinations around its initial point. Then, the network started training on the initial data and making predictions for the voltages that maximise fluorescence. Since the ion trap setup is very sensitive to changes in the electric field, the voltages were set to move a maximum of 1% of their previous value in each iteration to reduce the chance of losing the ion. As a step size value could not be explicitly defined, this percentage was chosen to make the changes similar to the step size used for ADAM.

A small percentage of our initial trials with the maximum change of a few percent (instead of 1%) led to an unstable ion during the parameter search sequence. This is because MLOOP is a global optimizer and can set the voltages to values far from the stable starting point. Since the ion trap is a complicated system that can only be modelled for a specific range of configurations, moving away from these settings can lead to unpredictable and usually unstable behavior. MLOOP also has an in-built mechanism that handles function noise using a predefined expected uncertainty. We set this uncertainty to the peak-to-peak noise of the photon readout when no optimization was running.

Since MLOOP is a global optimizer it was able to find optimum points different from the points found by ADAM. For trials where low numbers of initial training data points were used, these configurations proved to be unstable and in most cases resulted in the loss of the ion. Unstable states were also observed occasionally if the optimizer was run for too long. With moderate-size training sets, MLOOP was able to find voltage settings with fluorescence rates similar or higher than optimum points found by ADAM as shown in Fig. 4a. Considering the long duration of the MLOOP iteration sequence and the possibility of finding unstable settings in volatile conditions, the test of optimization with induced changing stray fields (“Testing under poor trap conditions”) was only performed with the ADAM optimizer as the gradient based search method proved to be more robust against fluctuations in the ion environment.

To test the effectiveness of the protocols, the saturation power, P_{sat} , was measured before and after the optimization process. The P_{sat} is the laser power at which the fluorescence rate of a two-level system is half the fluorescence at infinite laser power. We also measured the overall detection efficiency η , the fraction of emitted photons which resulted in detection events. Table 1 shows P_{sat} decreased (ion photon absorption was improved) using both ADAM and MLOOP. The detection efficiency was approximately the same for all runs as expected.

Another test was done by measuring fluorescence versus laser detuning before and after optimization. Figure 4b shows that the measured values follows the expected Lorentzian profile^{29–31} and associated linewidth before and after optimization. This indicates that the initial micromotion magnitude β was sufficiently small for fluorescence to be a good optimization proxy. Clear increase in fluorescence can be seen after optimizing 44

	P_{sat} (μW)	η (%)	Bkg (1/s)
Initial	4.11(1)	0.971(2)	1184(34)
ADAM-45	1.86(1)	0.938(1)	1227(35)
MLOOP-45	1.86(1)	0.989(1)	1291(35)

Table 1. Saturation power (P_{sat}), detection efficiency (η) and average background count rate (Bkg) before and after optimization with different methods. P_{sat} and η values are calculated by fitting the theoretical formula for Fluorescence Rate versus Laser Power to the experimental data^{29–31}. A clear drop in saturation power is observed after ADAM and MLOOP optimization of individual electrodes indicating a reduction of the micromotion.

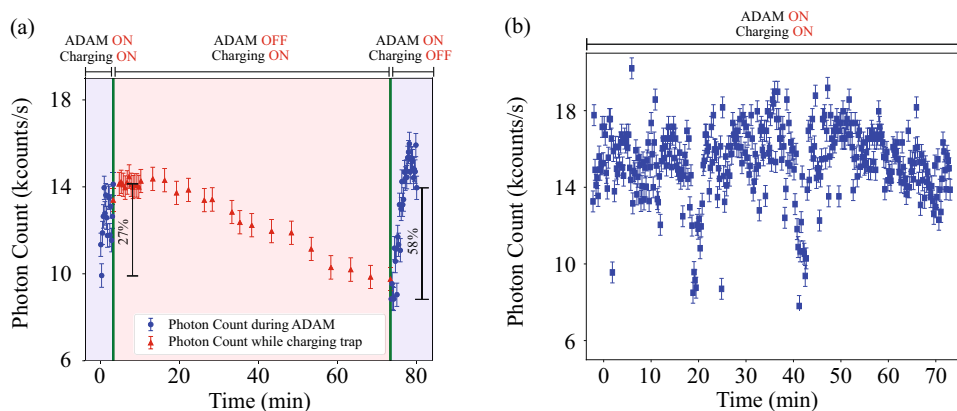


Figure 5. (a) Real time compensation with ADAM of laser charging induced stray electric field. The ion was optimized using ADAM (left blue points) then the photon count was noted whilst charging for 70 min (red points) then re-optimized (right blue points). Initial improvement from manually optimized settings was 27%. The second optimization improved the fluorescence by 58% from the charged conditions and returned it back to the optimum value of the first optimization within the error. (b) The trap was charged by hitting a UV laser to destabilize the ion and individual electrodes optimized using ADAM simultaneously for 70 min. The photon count fluctuates as a result of combination of fluctuations of power of the cooling laser power, algorithm search and charging irregularities. The optimizer keeps the fluorescence at the photon count similar to the case of optimizing after the charging is stopped (third section of (a)).

electrodes individually both with ADAM and MLOOP. The fit residual curve (difference between the experimental values and the theoretical fit) shows that optimizing individual electrodes, resulted in slight increase in heating instability near the resonance.

Testing under poor trap conditions. To test the live performance of the optimization protocol in a non-ideal situation, we deliberately charged the trap by shining 369.5 nm UV laser light onto the chip for 70 min. The power of the laser was $200 \pm 15 \mu W$ and the Gaussian diameter of the focus was $120 \pm 10 \mu m$. This process ejects electrons due to the photo-electric effect³⁴ and produces irregular and potentially unpredictable slow time varying electric fields within the trap. The process charged the trap significantly and made a noticeable reduction to the photon count. The ADAM algorithm was then tested both during charging and after charging was stopped. In both cases an improvement of fluorescence rate was observed.

The first experiment was performed to test the optimizing process after charging. In this test, starting with the optimal manual setting, ADAM individual electrode optimizer was able to obtain 27% improvement in the fluorescence rate (blue points on the left side of Fig. 5a). Then charging was induced onto the trap for 70 min and a clear decrease in photon count was seen that went even lower than the initial value (red points in Fig. 5a). At this point charging was stopped and ADAM was run again and fluorescence rate returned back to the previous optimum, within the error, in approximately 12 min. During the second optimization, the fluorescence goes higher than the stable final value for some iterations before the final. This is because of the same effect explained in “Gradient descent optimizer” section that the fluorescence might spike right after a change but go down slightly after stabilizing. Looking at the changes of individual electrodes, shown in Fig. 3b, we see that the main electrodes adjusted were those around the ion and some throughout the trap. The change in the laser horizontal position was negligible.

Another experiment was done by running ADAM during continuous charging for real-time compensation. Since we induce charging via laser scattering from the trap, the collected photons are both from the ion and the scattered laser and fluctuations in the intensity of scattered light confuses the optimizer. Despite that the optimizer did not lose the ion nor needed to abort the process. Figure 5b shows that the fluorescence rate, even after

a 70-min charging session, remained near the optimum value. After stopping the charging, the ion remained trapped for more than 8 h and was intentionally removed from the trap after this time.

Conclusion

Comparing the results of the gradient descent method and MLOOP shows that ADAM, the local optimizer, was a better performer in volatile conditions. This is because unlike MLOOP, ADAM does not need any training and can react faster to changes. MLOOP as a global optimizer was able to find settings with higher fluorescence but occasionally optimums were unstable if the defined parameter search range for the settings were large or if the machine learner ran for too long. The main drawback of MLOOP was that it was much slower than ADAM. In both cases photon collection times were the limiting factor.

Optimizing the individual electrodes improved the fluorescence by 78% for ADAM and 96% for MLOOP. ADAM showed its reliability in the charging test, returning the ion to the optimal fluorescence in 12 min after 70 min of purposely charging the trap with a UV laser. Improvement amounts depended greatly on the overall stability of the cooling laser and efficiency of the photon counting system.

To improve the speed of the gradient based optimizer, shorter readout times can be used although this increases variability in the value of the photon count. It is also possible to try local optimizers like SPSA³⁵ that approximate the gradient using only 2 readouts instead of 90. This improves the time needed for each iteration but because using this approximate gradient reduces the improvement in each iteration, the overall speed comparison to reach the same optimum can only be made when both are tried in practice. One possible problem with SPSA is that it cannot handle noise like ADAM and multiple readouts and averaging are needed in noisy conditions. If high amount of averaging is required, the algorithm may become very slow.

Received: 14 January 2022; Accepted: 8 April 2022

Published online: 29 April 2022

References

1. Debnath, S. *et al.* Demonstration of a small programmable quantum computer with atomic qubits. *Nature* **536**, 63–66. <https://doi.org/10.1038/nature18648> (2016).
2. Lekitsch, B. *et al.* Blueprint for a microwave trapped ion quantum computer. *Sci. Adv.* **3**, e1601540. <https://doi.org/10.1126/sciadv.1601540> (2017).
3. Monroe, C. & Kim, J. Scaling the ion trap quantum processor. *Science* **339**, 1164–1169. <https://doi.org/10.1126/science.1231298> (2013).
4. Blinov, B. B., Leibfried, D., Monroe, C. & Wineland, D. J. Quantum computing with trapped ion hyperfine qubits. *Quantum Inf. Process.* **3**, 45–59. <https://doi.org/10.1007/s11128-004-9417-3> (2004).
5. Wineland, D. J. *et al.* Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *J. Res. Natl. Inst. Stand. Technol.* **103**, 259–328. <https://doi.org/10.6028/jres.103.019> (1998).
6. Kielpinski, D., Monroe, C. & Wineland, D. J. Architecture for a large-scale ion-trap quantum computer. *Nature* **417**, 709–711. <https://doi.org/10.1038/nature00784> (2002).
7. Chiaverini, J. *et al.* Surface-electrode architecture for ion-trap quantum information processing. *Quantum Inf. Comput.* **5**, 419–439 (2005).
8. Seidelin, S. *et al.* Microfabricated surface-electrode ion trap for scalable quantum information processing. *Phys. Rev. Lett.* **96**, 253003. <https://doi.org/10.1103/PhysRevLett.96.253003> (2006).
9. Doret, S. C. *et al.* Controlling trapping potentials and stray electric fields in a microfabricated ion trap through design and compensation. *New J. Phys.* **14**, 073012. <https://doi.org/10.1088/1367-2630/14/7/073012> (2012).
10. Amini, J. M. *et al.* Toward scalable ion traps for quantum information processing. *New J. Phys.* **12**, 033031. <https://doi.org/10.1088/1367-2630/12/3/033031> (2010).
11. Kielpinski, D. Entanglement and decoherence in a trapped-ion quantum register. Ph.D. thesis (2001).
12. Tanaka, U. *et al.* Micromotion compensation in a surface electrode trap by parametric excitation of trapped ions. *Appl. Phys. B* **107**, 907–912. <https://doi.org/10.1007/s00340-011-4762-2> (2012).
13. Allcock, D. T. C. *et al.* Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap. *Appl. Phys. B* **107**, 913–919. <https://doi.org/10.1007/s00340-011-4788-5> (2012).
14. Narayanan, S. *et al.* Electric field compensation and sensing with a single ion in a planar trap. *J. Appl. Phys.* **110**, 114909. <https://doi.org/10.1063/1.3665647> (2011).
15. Berkeland, D. J., Miller, J. D., Bergquist, J. C., Itano, W. M. & Wineland, D. J. Minimization of ion micromotion in a Paul trap. *J. Appl. Phys.* **83**, 5025–5033. <https://doi.org/10.1063/1.367318> (1998).
16. Noek, R. *et al.* Trapping and cooling of $^{174}\text{Yb}^+$ ions in a microfabricated surface trap. *J. Korean Phys. Soc.* **63**, 907–913. <https://doi.org/10.3938/jkps.63.907> (2013).
17. Allcock, D. T. C. *et al.* Implementation of a symmetric surface-electrode ion trap with field compensation using a modulated Raman effect. *New J. Phys.* **12**, 053026. <https://doi.org/10.1088/1367-2630/12/5/053026> (2010).
18. Harlander, M., Brownnutt, M., Hänsel, W. & Blatt, R. Trapped-ion probing of light-induced charging effects on dielectrics. *New J. Phys.* **12**, 093035. <https://doi.org/10.1088/1367-2630/12/9/093035> (2010). [arXiv:1004.4842](https://arxiv.org/abs/1004.4842).
19. Seif, A. *et al.* Machine learning assisted readout of trapped-ion qubits. *J. Phys. B At. Mol. Opt. Phys.* **51**, 174006. <https://doi.org/10.1088/1361-6455/aad62b> (2018).
20. Wigley, P. B. *et al.* Fast machine-learning online optimization of ultra-cold-atom experiments. *Sci. Rep.* **6**, 25890. <https://doi.org/10.1038/srep25890> (2016).
21. Tranter, A. D. *et al.* Multiparameter optimisation of a magneto-optical trap using deep learning. *Nat. Commun.* **9**, 4360. <https://doi.org/10.1038/s41467-018-06847-1> (2018).
22. Ghadimi, M. *et al.* Scalable ion-photon quantum interface based on integrated diffractive mirrors. *npj Quantum Inf.* **3**, 1–4. <https://doi.org/10.1038/s41534-017-0006-6> (2017).
23. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. (2017). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
24. Shappert, C. M. *et al.* Spatially uniform single-qubit gate operations with near-field microwaves and composite pulse compensation. *New J. Phys.* **15**, 083053. <https://doi.org/10.1088/1367-2630/15/8/083053> (2013).
25. Blüms, V. *et al.* Laser stabilization to neutral Yb in a discharge with polarization-enhanced frequency modulation spectroscopy. *Rev. Sci. Instrum.* **91**, 123002. <https://doi.org/10.1063/5.0019252> (2020).
26. Norton, B. G., Streed, E. W., Petrasianas, M. J., Jechow, A. & Kielpinski, D. Millikelvin spatial thermometry of trapped ions. *New J. Phys.* **13**, 113022. <https://doi.org/10.1088/1367-2630/13/11/113022> (2011).

27. Budker, D., Kimball, D., Kimball, D. F. & DeMille, D. P. *Atomic Physics: An Exploration Through Problems and Solutions* (Oxford University Press, Oxford, 2004).
28. Zappacosta, A. Auto ion compensation. https://github.com/alexzappa11/Auto_Ion_Compensation (2021).
29. Metcalf, H. J., Van der Straten, P. & Van der Straten, P. *Laser Cooling and Trapping. Graduate Texts in Contemporary Physics* (Springer-Verlag, New York, 1999).
30. Ejtemaee, S., Thomas, R. & Haljan, P. C. Optimization of yb171+ fluorescence and hyperfine-qubit detection. *Phys. Rev. A* **82**, 063419. <https://doi.org/10.1103/PhysRevA.82.063419> (2010).
31. Berkeland, D. J. & Boshier, M. G. Destabilization of dark states and optical spectroscopy in Zeeman-degenerate atomic systems. *Phys. Rev. A* **65**, 033413. <https://doi.org/10.1103/PhysRevA.65.033413> (2002).
32. Ghadimi, M. *et al.* Multichannel optomechanical switch and locking system for wavemeters. *Appl. Opt.* **59**, 5136–5141. <https://doi.org/10.1364/AO.390881> (2020).
33. Price, K., Storn, R. M. & Lampinen, J. A. Natural computing series. In *Differential Evolution: A Practical Approach to Global Optimization* (Springer-Verlag, Berlin Heidelberg, 2005).
34. Wang, S. X. *et al.* Laser-induced charging of microfabricated ion traps. *J. Appl. Phys.* **110**, 104901, <https://doi.org/10.1063/1.3662118> (2011). [arXiv:1108.0092](https://arxiv.org/abs/1108.0092).
35. Spall, J. C. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Trans. Aerosp. Electron. Syst.* **34**, 817–823. <https://doi.org/10.1109/7.705889> (1998).

Acknowledgements

This research was financially supported by the Griffith University Research Infrastructure Programme, the Griffith Sciences equipment scheme, and Australian Research Council Linkage (LP180100096) Projects, ML is supported by Australian Research Council Future Fellowship (FT180100055), AZ was supported by Centre for Quantum Dynamics, Griffith University, Summer Scholarship, JS, and KS were supported by the Australian Government Research Training Program Scholarship.

Author contributions

A.Z. and M.G. developed the codes. A.Z., M.G., J.S., and K.S. performed the experiments. M.G., E.S. and M.L. supervised the research. A.Z., M.G. and M.L. prepared the figures. M.G., A.Z., E.S. and M.L. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.G.

Reprints and permissions information is available at 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 licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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 licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022