

SCIENTIFIC REPORTS



OPEN

Fast machine-learning online optimization of ultra-cold-atom experiments

Received: 21 December 2015

Accepted: 21 April 2016

Published: 16 May 2016

P. B. Wigley¹, P. J. Everitt¹, A. van den Hengel², J. W. Bastian³, M. A. Sooriyabandara¹, G. D. McDonald¹, K. S. Hardman¹, C. D. Quinlivan¹, P. Manju¹, C. C. N. Kuhn¹, I. R. Petersen⁴, A. N. Luiten⁵, J. J. Hope⁶, N. P. Robins¹ & M. R. Hush⁴

We apply an online optimization process based on machine learning to the production of Bose-Einstein condensates (BEC). BEC is typically created with an exponential evaporation ramp that is optimal for ergodic dynamics with two-body s-wave interactions and no other loss rates, but likely sub-optimal for real experiments. Through repeated machine-controlled scientific experimentation and observations our ‘learner’ discovers an optimal evaporation ramp for BEC production. In contrast to previous work, our learner uses a Gaussian process to develop a statistical model of the relationship between the parameters it controls and the quality of the BEC produced. We demonstrate that the Gaussian process machine learner is able to discover a ramp that produces high quality BECs in 10 times fewer iterations than a previously used online optimization technique. Furthermore, we show the internal model developed can be used to determine which parameters are essential in BEC creation and which are unimportant, providing insight into the optimization process of the system.

Experimental research into quantum phenomena often requires the optimization of resources or processes in the face of complex underlying dynamics and shifting environments. For example, creating large Bose-Einstein condensates (BECs) with short duty cycles is one of the keys to improving the sensitivity of cold-atom based sensors¹ or for performing scientific investigation into condensed matter phases², many-body physics³ and non-equilibrium dynamics⁴. The standard process of BEC production is evaporative cooling⁵. Microscopic semi-classical theory exists to describe this process⁶, but it can oversimplify the dynamics and miss more complex and effective methods of performing evaporation. For example, Shobu *et al.*⁷ found circumventing higher order inelastic collisions can produce large condensates. ‘Tricks’ like this are likely to exist for other species with complicated scattering processes⁸, but discovery is only possible by experimentation. We automate this process of discovery with *machine-learning* online optimization (MLOO). What distinguishes our approach from previous methods for optimization is that we seek to develop a statistical model of the relationship between parameters and the outcome of the experiment. We demonstrate that MLOO can discover condensation with less experiments than a competing optimization method and provide insight into which parameters are important in achieving condensation.

Online optimization (OO), with mostly genetic^{9–22} but also gradient²³ and hybrid solvers^{24,25}, has been used to enhance a variety of quantum experiments. Here we use online to mean optimization that is performed in real time with the experiment. What distinguishes our approach is that it does not only seek to optimize the experiment, but also creates an internal model that is able to predict the performance of future experiments given any set of parameters. This is achieved by modeling the experiment using a Gaussian process (GP)²⁶. Our algorithm both fits the model to previous observations, and chooses to do future experiments that will best refine its model, making it an automation of scientific method. Online machine learning (OML) with GPs^{26–30} has been applied in

¹Quantum Sensors and Atomlaser Lab, Department of Quantum Science, Research School of Physics and Engineering, The Australian National University, Acton, 2601, Australia. ²Australian Centre for Visual Technologies, University of Adelaide, Adelaide, 5005, Australia. ³School of Computer Science, University of Adelaide, Adelaide, 5005, Australia. ⁴School of Engineering and Information Technology, University of New South Wales at the Australian Defence Force Academy, Canberra, 2600, Australia. ⁵Institute for Photonics & Advanced Sensing, School of Physical Sciences, The University of Adelaide, Adelaide, 5005, Australia. ⁶Department of Quantum Science, Australian National University, Canberra, 2601, Australia. Correspondence and requests for materials should be addressed to M.R.H. (email: M.Hush@adfa.edu.au)




Figure 1. The experiment and the ‘learner’ form a closed loop. The learner produces a parameter set, X , for the experiment to test, these are converted into cooling ramps and used to perform an experiment. After the evaporation process is finished, an image of the cold atoms taken is used to calculate a cost function based on its quality as a resource $C(X)$. $C(X)$ is then fed back to the learner. The learner uses a GP to model the relationship between the input parameter values and the cost function values produced by the experiment. This model depends on a set of correlation lengths, or hyperparameters. Part (a) of the figure plots a set of observed costs (black circles with bars for uncertainty) with three possible GP models fit to the data: one with a long correlation length (red dotted), a medium correlation (blue solid) and a short correlation length (green dashed). Each GP is illustrated by a mean cost function bracketed by two curves indicating the function \pm one standard deviation from the mean. The correlation lengths affect both the mean and uncertainty of the model; note that the uncertainty approaches zero near the observed points. A final cost function is produced as a weighted average over the correlation lengths. This model is used to pick the next parameters X for the experiment.

a variety of areas including robotics^{31,32}, vision³³, industrial chemistry^{34,35} and biochemistry³⁶. However, in all of these cases, the focus was not on optimization. Rather, the goal was the development of an accurate model. We combine the advantages of OML with the motivation of OO. The resultant MLOO algorithm has the following advantages: every experimental observation is used to improve the GP model, and uncertainties in the measurements are correctly accounted for; our algorithm will find global minima, but exploration is not random, new parameters are picked with knowledge of where the learner is most uncertain; the learner provides a visualization of the resource’s quality as function of the parameters that can inform experimentalists on how to best develop future optimization experiments.

Methods

Experiment. The experimental apparatus is described in detail in³⁷. Initially ⁸⁷Rb atoms are cooled in a combined 2D and 3D MOT system and subsequently cooled further by RF (radio frequency) evaporation. The cloud is then loaded into a cross beam optical dipole trap for the final evaporation stage. It is this stage that is the subject of the optimization process. At this point, the sample contains 4×10^7 atoms at a temperature of $\sim 5 \mu\text{K}$ with a phase space density of ~ 0.05 . The cross dipole trap is formed from two intersecting 1090 nm and 1064 nm lasers with approximate waists of $350 \mu\text{m}$ and $300 \mu\text{m}$ respectively producing a trap with frequencies $185 \times 185 \times 40 \text{ Hz}$. The depth of the cross trap is determined by the intensity of the two beams and is found to be approximately $70 \mu\text{K}$. The 1064 nm beam is controlled by varying the current to the laser, while the 1090 nm beam is controlled using the current and a waveplate rotation stage combined with a polarizing beamsplitter to provide additional power attenuation while maintaining mode stability. A diagram of the experimental set up is shown in Fig. 1. Normally the power to these beams is ramped down over time, thereby lowering the walls of the trap and allowing the higher energy atoms to leak out. The remaining atoms rethermalize to a lower temperature, enabling cooling. Once the gas has been cooled to temperatures on the order of nK, a phase transition occurs, and a macroscopic number of atoms start to occupy the same quantum state. This transition is called Bose-Einstein condensation³⁸. We hand over control of these ramps to the MLOO. We consider two parameterizations: one simple, where we only control the start and end points of a linear interpolation; and one complex, where we add variable quadratic, cubic and quartic corrections to the simple case (see Supplemental Equations).

Performance Measure. The approach we propose is a form of supervised learning, meaning that we provide the learner with a number that quantifies the quality of the resource produced or in optimization terminology a cost that must be minimized. Naïvely one might try to use a measure based on temperature and particle number. However determining these quantities accurately near condensation is difficult when constrained to very few runs per parameter set. Instead, a technique was created to measure the width of the edges of the cloud. For thermal clouds this edge is broad, but as the sample cools and condenses these edges become sharper. To quantify this, an absorption image of the final state of the quantum gas is taken after a 30 ms expansion of the cloud, with the image providing the optical depth as a function of space. This absorption image is taken at resonance, resulting in saturation of the image (see Fig. 2). Whilst this makes determining peak density difficult, it ensures that the edges of the cloud are accurately determined. The cost is then calculated from all data between a lower and upper threshold optical depth. The lower threshold is determined by the noise in the system. The upper threshold is set slightly lower than the saturation level of the image. Only data from between the bounds is used and the cost is simply the average of these values. In practice this means the sharper the edges of the cloud, the lower the cost. Indeed, low quality thermal clouds have broad edges, whereas the ideal BEC has much sharper edges. Each parameter set is tested twice with the average of the two runs used for the cost. Tests of the variation in cost for a set of parameters run-to-run indicate they obey a Gaussian distribution. As such we are able to estimate the uncertainty from two

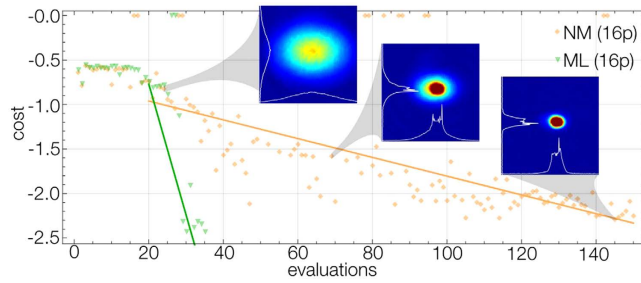


Figure 2. The optimization of the evaporation stage of creating a BEC using the complex 16 parameter scheme. The first 20 evaluations are an initial training run using a simple Nelder-Mead algorithm. The machine learning algorithm (green) then quickly optimizes to BEC. The insets show the different regimes experienced by the experiment, from a completely Gaussian thermal distribution, through the bimodal distribution containing a thermal background to the sharp edged BEC. The included cross-section illustrates how the cost decreases as the edges of the cloud get sharper.

runs as twice the range. In doing so, the chance we have underestimated the uncertainty will be 27%. We therefore also apply bounds to the uncertainty to eliminate outliers overly affecting the modeling process. The cost function can be evaluated as long as some atoms are present at the end of the evaporation run. In cases where the evaporation parameters produced no cloud twice for a set of parameters, we set the cost to a default high value.

Algorithm. We treat the experiment as a stochastic process $\mathcal{C}(X)$ which is dependent on the parameters $X = (x_1, \dots, x_M)$. When we make a measurement and determine a cost, we interpret this as a sample of this process $\mathcal{C}(X)$ with some associated uncertainty $U(X)$. We define the set of all parameters, costs and uncertainty previously measured as $\mathcal{X} = (X_1, \dots, X_N)$, $\mathcal{C} = (C_1, \dots, C_N)$ and $\mathcal{U} = (U_1, \dots, U_N)$ respectively and collectively refer to these sets as our observations $\mathcal{O} = (\mathcal{X}, \mathcal{C}, \mathcal{U})$. The aim of OO is to use previous observations \mathcal{O} to plan future experiments in order to find a set of parameters that minimize the mean cost of the stochastic process $M_{\mathcal{C}}(X)$. Unique to the MLOO approach, we first make an estimate of the stochastic process given our observations $\hat{\mathcal{C}}(X|\mathcal{O})$, which is then used to determine what parameters to try next.

We model $\mathcal{C}(X)$ as a GP—a distribution over *functions*—with constant mean function and covariance defined by a squared exponential correlation function $K(X, X', H) = e^{-\sum_{j=1}^M (x_j - x'_j)^2 / h_j^2}$ where $H = (h_1, \dots, h_M)$ is a set of correlation lengths for each of the parameters. The mean function conditional on the observations \mathcal{O} and correlation lengths H of our GP is: $\mu_{\hat{\mathcal{C}}}(X|\mathcal{O}, H)$, which is evaluated through a set of matrix operations²⁶ (see Supplemental Equations). As we are using a GP, we can also get the variance of the functions conditioned on \mathcal{O} and H : $\sigma_{\hat{\mathcal{C}}}(X|\mathcal{O}, H)$ ²⁶. Both of these estimates depend on the correlation lengths H , normally referred to as the hyperparameters of our estimate. We assume that H is not known a priori and needs to be fitted online.

The correlation lengths H control the sensitivity of the model to each of the parameters, and relates to how much a parameter needs to be changed before it has a significant effect on the cost (see Fig. 1). A standard approach to fit H is maximum likelihood estimation²⁶. Here, the hyperparameters are globally optimized over the likelihood of the parameters H given our observations \mathcal{O} , or $L(H|\mathcal{O})$ ²⁶ (see Supplemental Equations). However, when the data set is small there will often be multiple local optima for the hyperparameters whose likelihoods are comparable to the maximum. We term these hyperparameters the hypothesis set $\mathcal{H} = (H_1, \dots, H_p)$ with corresponding likelihood set $\mathcal{L} = (L_1, \dots, L_p)$.

To produce our final estimates for the mean function and variance we treat each hypothesis as a *particle*³⁰, and perform a weighted average over \mathcal{H} . The weighted mean function is now defined as $M_{\hat{\mathcal{C}}}(X|\mathcal{O}, \mathcal{H}) \equiv \sum_{i=1}^p w_i \mu_{\hat{\mathcal{C}}}(X|\mathcal{O}, \mathcal{H}_i)$ and weighted variance of the functions is $\Sigma_{\hat{\mathcal{C}}}^2(X|\mathcal{O}, \mathcal{H}) \equiv \sum_{i=1}^p w_i (\sigma_{\hat{\mathcal{C}}}^2(X|\mathcal{O}, \mathcal{H}_i) + \mu_{\hat{\mathcal{C}}}^2(X|\mathcal{O}, \mathcal{H}_i) - M_{\hat{\mathcal{C}}}^2(X|\mathcal{O}, \mathcal{H}))$, where $w_i = L_i / \sum_{i=1}^p L_i$ are the relative weights for the hyperparameters. Now that we have our final estimate for $\mathcal{C}(X|\mathcal{O}, \mathcal{H})$, we need to determine an optimization strategy for picking the next set of parameters to test.

Consider the following two strategies: We could always test parameters that are predicted to minimize $M_{\hat{\mathcal{C}}}(X)$, making our learner act as an ‘optimizer’. But this learner could get trapped in local minima and re-enforce its ignorance; Or we could test parameters that maximize $\Sigma_{\hat{\mathcal{C}}}^2(X)$ (i.e. where we are most uncertain), this would provide us with experimental data that helps us best refine our model and discriminate between the hypotheses, making our learner act like a ‘scientist’. But this learner may require a large number of trials to map the space and would not prioritize refinement of the global minima. We chose to implement a balanced strategy that repeatedly sweeps between these two extremes by minimizing a biased cost function: $B_{\hat{\mathcal{C}}}(X) \equiv bM_{\hat{\mathcal{C}}}(X) - (1 - b)\Sigma_{\hat{\mathcal{C}}}(X)$, where the value for b is linearly increased from 0 to 1 in a cycle of length Q . During testing with synthetic data, we found sweeping the learner between acting like a ‘scientist’ ($b = 0$) and an ‘optimizer’ ($b = 1$) was more robust and efficient than fixing the learner to one strategy. When we minimized $B_{\hat{\mathcal{C}}}(X)$ we also put bounds, set to 20% of the parameters maximum-minimum values, on the search relative to the last best measured X . We call these bounds a leash, as it restricts how fast the learner could change the parameters but did not stop it from exploring the full space (similar to trust-regions^{39,40}). This was a technical requirement for our experiment: when a set of




Figure 3. Optimization of evaporation curves to produce a BEC. The first 2N evaluations use a simple Nelder-Mead algorithm to learn about the cost space. The machine learning algorithm (red and blue) optimizes to BEC faster than the Nelder-Mead (black). By utilizing the machine learning model a parameter is eliminated and the convergence improves (red).

parameters was tested that was very different from the last set, the experiment almost always produced no atoms, meaning we had to assign a default cost that did not provide meaningful gradient information to the learner. Once the next set of parameters is determined they are sent to the experiment to be tested. After the resultant cost is measured this is then added to the observation set \mathcal{O} with $N \rightarrow N + 1$ and the entire process is repeated.

As a benchmark for comparison, we also performed OO using a Nelder-Mead solver⁴¹, which has previously been used to optimize quantum gates²⁵.

Results

We demonstrate the performance of machine learning online optimization in comparison to the Nelder-Mead optimizer in Fig. 2. Here we used the complex parameterization for all 3 ramps, and added an extra parameter that controlled the total time of the ramps, resulting in 16 parameters. If we were to perform a brute force search and optimize the parameters to within a 10% accuracy of the parameters maximum-minimum bounds, the number of runs required would be 10^{16} . The Nelder-Mead algorithm is able to find BEC much faster than this, in only 145 runs. The machine learning algorithm, on the other hand, is much faster. After the first 20 training runs, where the machine learning and Nelder-Mead algorithm use a common set of parameters, the machine learning algorithm converges in only 10 experiments.

The learner used in Fig. 2 only used the best hypothesis set when picking the next parameters, in other words we set $P = 1$. Evaluating multiple GPs is computationally expensive with so many parameters, so to save time we made this restriction. In spite of this, the learner discovered ramps that produced BEC in very few iterations. This is because the learner consistently fitted the correlation lengths of the 3 most important parameters—the end points of the ramps—very quickly. However, we found the other correlation lengths were not estimated well and would not converge, even after a BEC was found. This meant that we were unable to make useful predictions about the cost landscape and we could not reliably determine what parameters were least important. The final optimized parameters produced a condensate with 5×10^5 atoms.

Gramacy *et al.*³⁰ have suggested that making good online estimation of the GP correlation lengths requires multiple particles. We considered achieving this goal in a different experiment as shown in Fig. 3. Here we used a learner with many particles $P = 16$, but had to use the simple parameterization for the ramps to save computational time. This resulted in a total of 7 parameters. We can see again the overall trend for the machine learner is still faster than Nelder-Mead, but less pronounced. More carefully estimating the correlation lengths has hindered the convergence rate compared to the 16 parameter case. Nevertheless, as we now have a more reliable estimate of the correlation lengths we can take advantage of a different feature of the learner.

In Fig. 4(a) we show estimates of the cost landscape as 1D cross sections about the best measured point. We plot the two most sensitive parameters and the least. We can see the least sensitive parameter appears to have no effect on the production of a BEC. This parameter corresponds to an intentionally added 7th parameter of the system that controls nothing in the experiment. Figure 4(a) shows the learner successfully identified this, even with such a small data set. After making this observation we can then reconsider the design of the optimization process and eliminate this parameter from the experiment.

In Fig. 3 we plot the machine learner optimization run with $P = 16$ but now with only 6 parameters. We can see the learner converges more rapidly than the 7 parameter case, and even produces a higher quality BEC. As the learner no longer takes extra runs to determine the importance of the useless 7th parameter, it achieves BEC rapidly. We plot a 2D cross section of the landscape against the two most sensitive parameters in Fig. 4(b) generated from the 6 parameter machine learning run. We can see there is a very sharp transition to BEC, as it exists in a very deep valley of the landscape.

The optimum values for each parameter of the 16 parameter MLOO run are shown in a table in the supplementary material. Plots of the optimal ramps for each of the five optimization runs discussed are also shown and display cases where the optimum ramp is non-monotonically decreasing. The experimental controls adjust the shape of the trapping beams in a non-trivial manner, so this does not ensure that the trap depth experienced by the BEC was also non-monotonically decreasing. A key strength of the MLOO process is that we were able to find the ramps for the control that maximized the BEC without performing a lengthy characterization of relationship between the controls and the potential. Instead, the MLOO directly characterizes the relationship between the




Figure 4. Plots of the cross sections through the minima of the cost landscape as predicted by the learner. In (a) the predicted cost is shown as a function of the end of the polarization ramp (red), the end of the dipole beam ramp (green) and the unconnected parameter (blue). The learner correctly identifies that the unconnected parameter does not have a significant effect on the production of BEC. In (b) a cross section of the 2 most sensitive parameters are plotted against cost.

controls and the quality of the outcome. This significantly decreases the system identification and analysis overhead when optimizing an experiment.

The topology of optimization landscapes has been suggested as an explanation for the profound discrepancy between the number of experiments required to do a brute force search, and the number of experiments required in practice when using OO. Specifically, in laser-aided quantum chemistry, under the assumption of controllability it has been proven that landscapes are ‘trap-free’⁴² (there has been further refinement^{43–47} and debate^{48–50} on the generality of the result). In our experiment we observed Nelder-Mead always found condensation, albeit slower than MLOO, even though it is a local solver susceptible to being trapped. This suggests our landscape is also ‘trap-free’, and perhaps there is a universal principle for all quantum OO systems.

The MLOO algorithm we developed is available online⁵¹ (it uses⁵² to evaluate the GPs); it can be immediately applied to experiments that have previously used OO: quantum chemistry⁹, femtosecond physics¹³, and quantum computing²⁵. Indeed, any automated experiment with a resource of measurable quality can be enhanced using MLOO.

References

- Robins, N. P., Altin, P. A., Debs, J. E. & Close, J. D. Atom lasers: Production, properties and prospects for precision inertial measurement. *Phys. Rep.* **529**, 265–296 (2013).
- Bloch, I., Dalibard, J. & Nascimbène, S. Quantum simulations with ultracold quantum gases. *Nat. Phys.* **8**, 267–276 (2012).
- Bloch, I., Dalibard, J. & Zwierger, W. Many-body physics with ultracold gases. *Rev. Mod. Phys.* **80**, 885–964 (2008).
- Langen, T., Geiger, R. & Schmiedmayer, J. Ultracold Atoms Out of Equilibrium. *Annu. Rev. Condens. Matter Phys.* **6**, 201–217 (2015).
- Olson, A. J., Niffenegger, R. J. & Chen, Y. P. Optimizing the efficiency of evaporative cooling in optical dipole traps. *Phys. Rev. A* **87**, 053613 (2013).
- Sackett, C. A., Bradley, C. C. & Hulet, R. G. Optimization of evaporative cooling. *Phys. Rev. A* **55**, 3797–3801 (1997).
- Shobu, T., Yamaoka, H., Imai, H., Morinaga, A. & Yamashita, M. Optimized evaporative cooling for sodium Bose-Einstein condensation against three-body loss. *Phys. Rev. A* **84**, 033626 (2011).
- Altin, P. A. *et al.* Collapse and three-body loss in a 85Rb Bose-Einstein condensate. *Phys. Rev. A* **84**, 033632 (2011).
- Judson, R. S. & Rabitz, H. Teaching lasers to control molecules. *Phys. Rev. Lett.* **68**, 1500–1503 (1992).
- Warren, W. S., Rabitz, H. & Dahleh, M. Coherent Control of Quantum Dynamics: The Dream Is Alive. *Science* **259**, 1581–1589 (1993).
- Amstrup, B., Toth, G. J., Szabo, G., Rabitz, H. & Loerincz, A. Genetic Algorithm with Migration on Topology Conserving Maps for Optimal Control of Quantum Systems. *J. Phys. Chem. A* **99**, 5206–5213 (1995).
- Dods, J., Gruner, D. & Brumer, P. A genetic algorithm approach to fitting polyatomic spectra via geometry shifts. *Chem. Phys. Lett.* **261**, 612–619 (1996).
- Baumert, T., Brixner, T., Seyfried, V., Strehle, M. & Gerber, G. Femtosecond pulse shaping by an evolutionary algorithm with feedback. *Appl. Phys. B* **65**, 779–782 (1997).
- Pearson, B. J., White, J. L., Weinacht, T. C. & Bucksbaum, P. H. Coherent control using adaptive learning algorithms. *Phys. Rev. A* **63**, 063412 (2001).
- Zeidler, D., Frey, S., Kompa, K.-L. & Motzkus, M. Evolutionary algorithms and their application to optimal control studies. *Phys. Rev. A* **64**, 023420 (2001).
- Walmsley, I. & Rabitz, H. Quantum physics under control. *Phys. Today* **56**, 43–49 (2003).
- Rohringer, W. *et al.* Stochastic optimization of a cold atom experiment using a genetic algorithm. *Appl. Phys. Lett.* **93**, 264101 (2008).
- Tsubouchi, M. & Momose, T. Rovibrational wave-packet manipulation using shaped midinfrared femtosecond pulses toward quantum computation: Optimization of pulse shape by a genetic algorithm. *Phys. Rev. A* **77**, 052326 (2008).
- Rohringer, W., Fischer, D., Schmiedmayer, J., Trupke, M. & Schumm, T. Stochastic optimization of Bose-Einstein condensation using a genetic algorithm (2011). URL <http://cdn.intechopen.com/pdfs-wm/13865.pdf>. Date of access:16/03/2016.
- Heinze, G., Hubrich, C. & Halfmann, T. Stopped Light and Image Storage by Electromagnetically Induced Transparency up to the Regime of One Minute. *Phys. Rev. Lett.* **111**, 033601 (2013).
- Starkey, P. T. *et al.* A scripted control system for autonomous hardware-timed experiments. *Rev. Sci. Instrum.* **84**, 085111 (2013).
- Geisel, I. *et al.* Evolutionary optimization of an experimental apparatus. *Appl. Phys. Lett.* **102**, 214105 (2013).
- Roslund, J. & Rabitz, H. Gradient algorithm applied to laboratory quantum control. *Phys. Rev. A* **79**, 053417 (2009).

24. Egger, D. & Wilhelm, F. Adaptive Hybrid Optimal Quantum Control for Imprecisely Characterized Systems. *Phys. Rev. Lett.* **112**, 240503 (2014).
25. Kelly, J. *et al.* Optimal Quantum Control Using Randomized Benchmarking. *Phys. Rev. Lett.* **112**, 240504 (2014).
26. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (the MIT Press, 2006).
27. Seo, S., Wallat, M., Graepel, T. & Obermayer, K. Gaussian process regression. *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks* **3**, 241–246 (2000).
28. Csató, L. & Oppier, M. Sparse on-line gaussian processes. *Neural Comput.* **14**, 641–668 (2002).
29. Deisenroth, M. P., Rasmussen, C. E. & Peters, J. Gaussian process dynamic programming. *Neurocomputing* **72**, 1508–1524 (2009).
30. Gramacy, R. B. & Polson, N. G. Particle Learning of Gaussian Process Models for Sequential Design and Optimization. *J. Comp. Graph. Stat.* **20**, 102–118 (2011).
31. Nguyen-Tuong, D. & Peters, J. Local Gaussian process regression for real-time model-based robot control. *IEEE/RSJ International Conference on Intelligent Robots and Systems* **1**, 380–385 (2008).
32. Nguyen-Tuong, D., Peters, J. R. & Seeger, M. Local Gaussian Process Regression for Real Time Online Model Learning. In *Advances in Neural Information Processing Systems* **21**, 1193–1200 (2009).
33. Ranganathan, A., Yang, M.-H. & Ho, J. Online Sparse Gaussian Process Regression and Its Applications. *IEEE Trans. Image Process.* **20**, 391–404 (2011).
34. Yu, J. Online quality prediction of nonlinear and non-Gaussian chemical processes with shifting dynamics using finite mixture model based Gaussian process regression approach. *Chem. Eng. Sci.* **82**, 22–30 (2012).
35. Li, X., SU, H. & CHU, J. Multiple Model Soft Sensor Based on Affinity Propagation, Gaussian Process and Bayesian Committee Machine. *Chinese J. Chem. Eng.* **17**, 95–99 (2009).
36. Gao, P., Honkela, A., Rattray, M. & Lawrence, N. D. Gaussian process modelling of latent chemical species: applications to inferring transcription factor activities. *Bioinformatics* **24**, i70–i75 (2008).
37. Kuhn, C. C. N. *et al.* A Bose-condensed, simultaneous dual-species Mach-Zehnder atom interferometer. *New J. Phys.* **16**, 073035 (2014).
38. Anglin, J. R. & Ketterle, W. Bose-Einstein condensation of atomic gases. *Nature* **416**, 211–218 (2002).
39. Conn, A. R., Gould, N. I. M. & Toint, P. L. *Trust Region Methods* (Society for Industrial and Applied Mathematics, 2000).
40. Yuan, Y. A review of trust region algorithms for optimization. *ICIAM* **99**, 271–282 (2000).
41. Nelder, J. A. & Mead, R. A Simplex Method for Function Minimization. *Comput. J.* **7**, 308–313 (1965).
42. Rabitz, H. A., Hsieh, M. M. & Rosenthal, C. M. Quantum Optimally Controlled Transition Landscapes. *Science* **303**, 1998–2001 (2004).
43. Chakrabarti, R. & Rabitz, H. Quantum control landscapes. *Int. Rev. Phys. Chem.* **26**, 671–735 (2007).
44. Wu, R., Rabitz, H. & Hsieh, M. Characterization of the critical submanifolds in quantum ensemble control landscapes. *J. Phys. A* **41**, 015006 (2008).
45. Wu, R., Pechen, A., Rabitz, H., Hsieh, M. & Tsou, B. Control landscapes for observable preparation with open quantum systems. *J. Math. Phys.* **49**, 022108 (2008).
46. Hsieh, M., Wu, R., Rosenthal, C. & Rabitz, H. Topological and statistical properties of quantum control transition landscapes. *J. Phys. B* **41**, 074020 (2008).
47. Ho, T.-S., Dominy, J. & Rabitz, H. Landscape of unitary transformations in controlled quantum dynamics. *Phys. Rev. A* **79**, 013422 (2009).
48. Pechen, A. N. & Tannor, D. J. Are there Traps in Quantum Control Landscapes? *Phys. Rev. Lett.* **106**, 120402 (2011).
49. Rabitz, H., Ho, T.-S., Long, R., Wu, R. & Brif, C. Comment on “Are There Traps in Quantum Control Landscapes?”. *Phys. Rev. Lett.* **108**, 198901 (2012).
50. Pechen, A. N. & Tannor, D. J. Pechen and Tannor Reply. *Phys. Rev. Lett.* **108**, 198902 (2012).
51. Hush, M. M-loop: Machine-learning online optimization package (2016). URL <https://github.com/michaelhush/M-LOOP>. Date of access: 16/03/2016.
52. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).

Acknowledgements

M.R.H. acknowledges funding from an Australian Research Council (ARC) Discovery Project (project number DP140101779). J.J.H. acknowledges support of an ARC Future Fellowship (FT120100291). A.N.L. would like to thank the South Australian Government through the Premier’s Science and Research Fund for supporting this work. I.R.P. acknowledges support of an ARC Laureate Fellowship FL110100020.

Author Contributions

P.J.E., K.S.H., C.C.N.K., G.D.M., N.P.R. and P.B.W. contributed to design and construction of the experimental apparatus. P.J.E. and P.B.W. collected and analyzed data. J.W.B., P.J.E., J.J.H., M.R.H., I.R.P. and P.B.W. contributed to designing the machine learning approach. P.B.W. and M.R.H. wrote the manuscript. M.R.H. wrote the M-LOOP software. A.v.d.H., J.J.H., M.R.H., A.N.L. and N.P.R. contributed to the original conception of the experiment. N.P.R. guided direction of experimental approach. All authors contributed to discussions of experiment and results, and assisted in editing the manuscript.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Wigley, P. B. *et al.* Fast machine-learning online optimization of ultra-cold-atom experiments. *Sci. Rep.* **6**, 25890; doi: 10.1038/srep25890 (2016).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>