Lucas Marcelo de Sá Marques dos Santos

Machine learning optimization of a magneto-optical trap

São Carlos

12/2020

Lucas Marcelo de Sá Marques dos Santos

Machine learning optimization of a magneto-optical trap

Course conclusion work presented to the Undergraduate Program in Physics at Instituto de Física de São Carlos da Universidade de São Paulo, to obtain the degree of Bachelor in Physics.

Advisor: Dr. Gustavo Deczka Telles

São Carlos

12/2020

dos Santos, Lucas Marcelo de Sá Marques
Machine learning optimization of a magneto-optical trap / Lucas Marcelo de Sá Marques dos Santos; advisor Gustavo Deczka Telles -- São Carlos, 2020.
31 p.

Course conclusion work (Bachelor's degree - Undergraduate Program in Theoretical and Experimental Physics) -- Instituto de Física de São Carlos, Universidade de São Paulo - Brasil, 2020.

1.Atomic cooling. 2.Magneto-optical trap. 3.Machine learning. I. Telles, Gustavo Deczka, advisor. II. Machine learning optimization of magneto-optical trap.

Aos meus pais,
Elizabeth Leoncio de Sá Marques dos Santos
e Marcelo Marques dos Santos.

Acknowledgements

I would like to extend my enormous gratitude to my supervisor, Dr. Gustavo Telles, for all of the guidance, openness to new ideas and enlightening discussions about both Physics and life as a researcher he has offered me over the past few years.

Abstract

As an essential component of modern production techniques of Bose-Einstein condensates, the magneto-optical trap has an important role in determining the size of the condensate that can be obtained. Extensive work has been done over the years with the goals both of simplifying the construction of the trap and of acquiring a deeper understanding of the mechanisms at play in its interior, which has also allowed, to an extent, the prior determination of adequate choices of experimental parameters that lead to the production of larger and colder samples. However, although a simple description is enough to provide a broad understanding of the most relevant parameters - e.g., the cooling beam detuning and the magnetic field gradient - still other phenomena arise in the trap which are not taken into account by this model, such as collective effects and the presence of an additional laser beam. This absence of complete theoretical modeling precludes precise, a priori, optimization of the magneto-optical trap. As a result, optimization is typically started from a set of standard, generally adequate parameters, which are then fine-tuned manually, a process that is of limited precision and, most importantly, too time-consuming to be done often enough to keep the trap near optimal conditions at all times. Machine learning optimization methods may provide an alternative that is both precise and time-efficient. While previous work has applied machine learning to the optimization of a magneto-optical trap with a performance measure based on optical depth, we propose defining this performance measure from the trap loading curve. In order to properly substantiate our proposal, we go over much of the basic physics of the magneto-optical trap and its loading dynamics, establishing the connection between experimental parameters and our performance measure. We then introduce our custom implementation and the methods utilized for optimization, which resulted in a considerable increase of the final trap population for shorter times and starting from unoptimized parameters.

Keywords: Atomic cooling. Magneto-optical trap. Machine learning.

Resumo

Como uma componente essencial de técnicas modernas de produção de condensados de Bose-Eistein, a armadilha óptico-magnética tem um importante papel na determinação do tamanho do condensado que pode ser obtido. Extenso trabalho tem sido feito ao longo dos anos com os objetivos tanto de simplificar a construção da armadilha quanto de adquirir um entendimento mais profundo dos mecanismos em ação no seu interior, o que também tem permitido, até certo ponto, uma determinação prévia da escolha de parâmetros adequados para a produção de amostras maiores e mais frias. Todavia, embora uma descrição simples seja suficiente para prover um entendimento geral dos parâmetros mais relevantes - e.g., o *detuning* do feixe de resfriamento e o gradiente do campo magnético - restam ainda outros fenômenos que não são levados em conta por este modelo na armadilha, como efeitos coletivos e a presença de feixes adicionais. Essa falta de modelagem teórica completa impede, a princípio, uma otimização precisa da armadilha óptico-magnética. Como resultado, otimização tipicamente começa com um conjunto de parâmetros padrão adequados, os quais passam então por um ajuste fino manual, um processo que é de precisão limitada e, mais importantemente, lento demais para ser feito com frequência suficiente para manter a armadilha em condições próximas das ótimas a todo momento. Métodos de otimização por aprendizado de máquina podem oferecer uma alternativa que é tanto precisa quanto eficiente. Enquanto trabalhos anteriores têm aplicado aprendizado de máquina na otimização de uma armadilha magneto-óptica com uma medida de desempenho baseada na profundidade óptica, propomos aqui definir essa medida de desempenho a partir da curva de carregamento da armadilha. De forma a substanciar apropriadamente nossa proposta, passamos por boa parte da física básica da armadilha magneto-óptica e sua dinâmica de carregamento, estabelecendo a conexão entre os parâmetros experimentais e nossa medida de desempenho. Introduzimos então nossa própria implementação e os métodos utilizados para a otimização, os quais resultaram em um crescimento considerável da população final da armadilha para tempos mais curtos e a partir de parâmetros não-otimizados.

Palavras-chave: Resfriamento atômico. Armadilha magneto-óptica. Aprendizado de máquina.

Contents

1 Introduction

The magneto-optical trap (MOT) is today a staple of the field of cold atoms, an essential first stage in the production of Bose-Einstein condensates (BEC) that is able to efficiently cool and trap large samples of atoms. Its basic principle is laser cooling, a technique which relies on the overall cooling effect that results from transfers of momentum to an atom upon the scattering of photons. Although atomic traps based on this scattering force alone were proposed from the inception of laser cooling, Ashkin and Gordon[1] proved, in 1975, what was called an optical Earnshaw's theorem: an atom cannot be trapped in free space solely through scattering forces from a static configuration of laser beams. This result, however, was later shown by Pritchard *et al.*[2] to not always apply; scattering forces could in fact be used to construct traps for neutral atoms, as long as these forces were not dependent solely on light intensity. It was in this same work that which would become known as the magneto-optical trap was first proposed, where the inclusion of a magnetic field gradient in the setup would, due to the Zeeman effect, add a spatial dependence to the scattering force, turning it into an effective trapping, in addition to cooling, force.

The MOT was first realized by Raab *et al.*[3] in 1987, at which time it still required a pre-cooling stage; only in 1990 did Monroe *et al.*[4] elaborate the vapor-cell magneto-optical trap, which captured thermal atoms directly from a background gas in the same cell as the trap itself, with no need for preliminary cooling. Thanks to this development, the MOT was finally made into a relatively simple and accessible means of obtaining and studying cold gases, which allowed a great deal of work to be done on the technique thereafter. Many different ways of refining the process would be developed in the following years, with the intention both of producing denser and larger samples of cold atoms, and of studying the physical phenomena taking place in the trap. These phenomena included effects particular to the conditions inside an optical trap, such as the radiative forces of attenuation and re-absorption,[5] and light-assisted collisions,[6] which had to be taken into account in order to improve the apparatus as they limited both the trap density and the steady state population. One of the main objectives of these improvements was to achieve Bose-Einstein condensation, which was finally realized by applying evaporative cooling to a magnetic trap after initial cooling in a magneto-optical trap in 1995. This would establish the magneto-optical trap's place as the standard first step in the creation of the Bose-Einstein condensate, but the advent of BECs would also, unfortunately, draw most attention away from studying the MOT for its own rich and interesting properties.

The theoretical treatment of the scattering force in a MOT is usually kept to its most fundamental model, which describes cooling under ideal conditions along the axes of the quadrupole field in the

trap space. More detailed theoretical treatment has sought to account for the Gaussian profile of the laser beam and excitation of multiple hyperfine transitions for atoms located away from the axes, but while general trends may be well predicted, even these more detailed models are not able to make precise analytical predictions.[7] As a result, optimization of the MOT is usually done through a manual search of the parameter space starting from typical, unoptimized configurations. This process is evidently still imprecise, as well as very time-consuming, which makes regularly optimizing the MOT unpractical, as it is just a single stage in a long and complex sequence for the production of a Bose-Einstein condensate.

On the other hand, machine learning (ML) is a technique that is becoming more and more common for tackling complex systems lacking a precise theoretical description. Of special interest are online algorithms, which can be trained incrementally as new data is made available, in contrast to batch learning, where a model must be fitted to all data at once. Because each training step is fast, the online optimization can be done on the fly: this means that, in the context of an experiment such as the MOT, an online learning algorithm can be updated in real time with consecutive runs, without a significant delay between them. If the algorithm is then given control of the experimental setup itself, this would allow it to effectively search within a certain parameter range in order to find the optimal values, much in the same way as manual optimization is realized, but in much less time and also more efficiently. This makes ML a promising candidate for optimization of a MOT, allowing for precise fine-tuning of the experiment without the inefficiency of manual optimization, while avoiding the shortcomings of the analytical models.

In this work, we have studied the application of Gaussian process regression, artificial neural network and differential evolution methods through the Machine Learning Online Optimization Package (`M-LOOP` hereon) toolkit to the optimization of a magneto-optical trap for the production of larger samples of trapped cold ^{87}Rb atoms. To this end, we have proposed the use of the MOT loading curve as a performance measure that makes faster loading and larger samples desirable for an optimization algorithm. We start by introducing the magneto-optical trap and its basic principles in sec. 2, where we define the main mechanisms at play for the cooling, discuss non-ideal aspects of the trap and study its loading dynamics in order to define a good model for the loading curve. In sec. 3 we describe our implementation and the three types of optimization algorithms investigated. Sec. 4 briefly summarizes the experimental setup utilized. In sec. 5 we discuss the results of our implementation, and then present our concluding remarks in sec. 6, alongside suggestions for future studies.

2 The magneto-optical trap

2.1 Setting up a magneto-optical trap

In the context of laser cooling, the scattering of light by an atom is the process through which an atom absorbs a photon, undergoes a transition to an excited state, and then returns to its ground state through spontaneous emission. The interacting radiation must then of course be resonant to some atomic transition. When immersed in a laser beam, the atom may also interact with a second resonant photon while in the excited state and decay through stimulated emission; the reemitted photon then follows the same direction and polarization as the stimulated one, which is also the direction of the originally absorbed photon.

The problem of the scattering force can be first treated as the one-dimensional case of a two-level atom, with resonance frequency ω_0, immersed in a time-varying and linearly polarized classical electrical field of frequency ω_L. The field interaction is a perturbation over the atomic Hamiltonian, and Schrödinger's time-dependent equation - in addition to manual inclusion of the decaying effect of spontaneous emission - leads us to the set of coupled rate equations known as the optical Bloch equations. Now, since the decay rate of an excited atom is $\Gamma = 1/\tau$, τ the lifetime of the excited state, an atom with a probability p of being in the excited state has an effective decay rate of Γp. Because in a laser beam the atom can scatter as soon as it decays, this is also the scattering rate $\Gamma_s = \Gamma p$. In the steady state, the optical Bloch equations give us

$$\Gamma_s = \frac{\Gamma}{2}\frac{S}{1 + S + (2\delta/\Gamma)^2}, \;\; S = I/I_s \tag{2.1}$$

Where I is the intensity of the laser beam, the saturation intensity I_s is that at which Rabi oscillations and spontaneous decay happen at similar rates and $\delta = \omega_L - \omega_0$ is the laser frequency detuning.

Having defined a scattering rate, we can then define a scattering force. Conservation of momentum tells us that, upon absorbing a photon, the atom receives a momentum kick equal to the original momentum of that photon. Conservation of momentum also tells us that, upon emission, the atom gains another momentum kick, opposite to the momentum of the emitted photon. Because spontaneous emission is isotropic, its effect on atomic momentum averages to zero over time, so that scattering does produce an effective force on the atom. The scattering force $\mathbf{F}_s$ due to a laser beam is then simply $\mathbf{F}_s = \hbar\Gamma_s\mathbf{k}$.

However, we must be mindful of the fact that we are always dealing with moving atoms. Thermal

atoms captured in a MOT have velocities on the order of 100 m/s, so that the Doppler effect comes into play to a significant degree in shifting the radiation frequency. If the atom moves against the beam, the radiation is blue shifted, whilst if it moves with the beam, the radiation is redshifted, but in any case the beam will imprint on the atom a force that points in the way light is propagating. In order for an atom that is moving against a beam to see it in resonance, the beams in a MOT are redshifted. Because the Doppler shift is equivalent in practice to an extra detuning, we may simply add the Doppler shift ω_D to the detuning δ in eq. (2.1).

Now consider the following. For a given fixed beam, the scattering force always has the same direction. This means that a counterpropagating atom will be cooled until it is practically still along the beam axis, but will then acquire an acceleration in the opposite direction, being heated. For two counterpropagating redshifted beams however, once it begins to move in the opposite direction, the atom starts to perceive the "heating beam" as more redshifted than the other. The net force for a redshifted beam pair is then effectively a viscous damping force, which has led to the cooling technique called *optical molasses* (OM): a set of three pairs of counterpropagating, redshifted beams, with cooling occurring in all directions in the region where the beams intersect.

However, effective as a cooling apparatus it is, OM is still unable to provide trapping as per Earnshaw's optical theorem. Velocity in an OM has a lower limit called the Doppler limit,[8] which gives velocities on the order of cm/s. With velocities so low, the Doppler shift does not bring redshifted radiation sufficiently close to resonance, and as a consequence, cold atoms interact so weakly with the light that they are able to simply diffuse out of the trap. A MOT solves this issue by introducing a spatially-varying magnetic field, which pushes the atoms "back into resonance" with the radiation if they stray too far from the center of the molasses. This happens because of the Zeeman effect, which splits the atom's energy levels, lifting the degeneracy of the magnetic quantum number. This splitting also affects the polarization of light associated with the split transition, which a simple classical description can help us understand.

In this classical model, the atom is looked at as a tridimensional isotropic harmonic oscillator of frequency ω_0 immersed in a homogeneous magnetic field $\mathbf{B} = B\hat{\mathbf{z}}$. Solving its equation of motion as an eigenvalue problem gives us three solutions: along the z axis with frequency ω_0, anti-clockwise rotation around z on the xy plane with frequency $\omega_0 - \Omega_L$ and clockwise rotation around z with frequency $\omega_0 + \Omega_L$. The oscillating electron emits dipole radiation, which is polarized: for $\omega = \omega_0$, radiation is emitted in all directions but z, and is always linearly polarized. In the cases $\omega = \omega_0 \pm \Omega_L$, light is radiated in all directions, but polarization depends on viewing angle. For an observer on the xy plane, it is again linearly polarized, while for an observer on z away from the origin it is circularly

polarized, with handedness depending on both motion and whether the observer is in $z > 0$ or $z < 0$.

For the exact shift, we resort to perturbation theory, so that the Zeeman shift is a first order correction for the hyperfine structure perturbed by the field. We get, for the m_F sublevel of the F hyperfine level, the Zeeman energy correction $\Delta E_{m_F} = g_F \mu_B B m_F$, where g_F is the Landé g-factor for the atom with total angular momentum F.

In a MOT, two coils in an anti-Helmholtz configuration are used to induce a quadrupole field $\mathbf{B}$ centered in the center of the beam overlap region, where it goes to zero, while along each axis it is parallel to that axis and has constant gradient. Taking the z axis to be the coil axis, the m_F level of an atom at coordinate z is accordingly Zeeman shifted by $\Delta E_{m_F}(z) = g_F \mu_B \frac{dB_z}{dz} m_F z$.

Let us take our two-level atom to have $F = 0$ for the ground state and $F' = 1$ for the excited state. Then the Zeeman effect splits the excited state into three hyperfine states of distinct energies, $m_{F'} = -1, 0, +1$. For an atom at $z > 0$, these energies are below, equal to and above the original, respectively; for $z < 0$, we have the inverse situation. Because for a slow atom both beams are redshifted, both of them are much more likely to drive the downshifted transition. We should then like to have the $+z$ propagating beam driving $m_F = 0 \to m_{F'} = 1$, thus only pushing atoms in $z < 0$, and the $-z$ propagating beam driving $m_F = 0 \to m_{F'} = -1$, thus only pushing atoms in $z > 0$. In the MOT, this is accomplished by appropriately using circularly polarized light: in order to drive $m_F = 0 \to m_{F'} = 1$, the z propagating beam is σ^+-polarized, while the $-z$ propagating beam is σ^--polarized in order to drive $m_F = 0 \to m_{F'} = -1$.

The Zeeman shift in the frequency itself is simply the difference between the Zeeman shifts of the excited and ground states. For a general $m_F \to m_{F'}$ transition

$$\hbar \omega_Z = (g_{F'} m_{F'} - g_F m_F) \mu_B \frac{dB_z}{dz} z = \mu' \frac{dB_z}{dz} z \tag{2.2}$$

Where we have defined an effective magnetic momentum μ' for the transition. As with the Doppler shift, the Zeeman shift ω_Z is in practice a simple addition to the detuning, so the full force is

$$\mathbf{F}_{\text{MOT}} = \frac{\hbar \Gamma \mathbf{k}}{2} \left[\frac{S}{1 + S + \left[2 \left(\delta + \omega_D + \omega_Z \right) / \Gamma \right]^2} - \frac{S}{1 + S + \left[2 \left(\delta - \omega_D - \omega_Z \right) / \Gamma \right]^2} \right] \tag{2.3}$$

Eq. (2.3) makes it very clear how the balance between the frequency shifts δ, ω_D and ω_Z determines the effectiveness of cooling. By controlling the frequency of the cooling beams, we can directly set δ and indirectly affect ω_D, and by controlling the current on the coils that generate the B field, we can indirectly control ω_Z. These two parameters, cooling frequency and coil current, are therefore prime targets for optimization of the trap. Fig. 1 illustrates their effect.

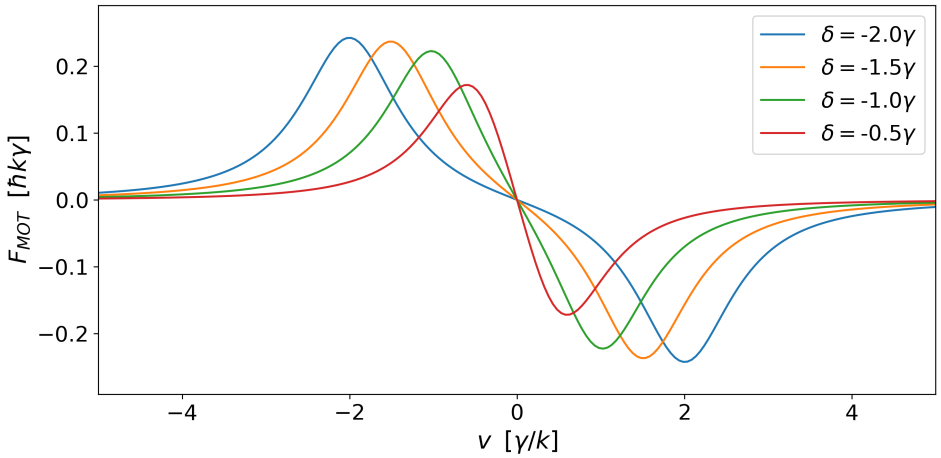

Figure 1: Scattering force relative to velocity for multiple detunings, at $z = 0$ and $S = 1$. The same behavior exists in relative to position for varying magnetic field gradients.
Source: By the author.

Although this model is developed for an ideal two-level atom, it describes well enough what happens in a real multilevel atom trapped in the MOT, provided that this atom is restricted to a single pair of ground and excited states. This is not the case, however, for the rubidium-87 atoms used in the MOT we have optimized. Albeit the probability of exciting a transition is maximal for resonant radiation, it does not immediately drop to zero as soon the radiation suffers a slight shift, and for ^{87}Rb the cooling beam still has a chance of exciting a particular off-resonant transition which eventually ”leaks” all atoms out of the cooling cycle. Our MOT then includes a repump beam responsible for exciting this ”repump transition” that returns leaked atoms to cooling. Our trap also includes an additional set of external (compensation) coils which play a role in further cooling stages and are not necessary for the MOT, but that have been observed to produce some improvement in the MOT when manually fine-tuned.

We then choose the repump frequency and the compensation coils current as a second pair of target parameters, alongside the cooling frequency and the main coils current. Through the repump frequency the number of atoms available for cooling may be controlled, while controlling the compensation current allows ML to be better compared with human optimization. Having defined our four optimization targets as thus, we now determine how we will provide our optimizer with a measure of the performance of a given set of parameters.

2.2 The loading curve

A good MOT is generally one that produces a large sample of cold atoms to be transported to further cooling stages of BEC production. Naturally, we would like to use the number of atoms loaded

to the trap as our performance measure for optimization. Full loading of a MOT typically takes 30 s to 40 s, which would be the time needed to collect a single training datum for an optimization algorithm, while, as we shall see, the best optimization may take close to a hundred points of data. Fortunately, it turns out that we don't actually need to fully load the trap for this.

In a vapor-cell MOT, atoms are captured directly from a background gas established inside the chamber where trapping occurs. But, as atoms are not invisible to each other, collisions provide a means for exchanges of energy that allow some atoms to escape from the trap. We begin by stating the empirical model most commonly used for a complete description of the loading curve. It is given in the form of a rate equation for the population N of the trap:[6]

$$\frac{dN}{dt} = R - \gamma N - \beta \int_V dV n^2 \tag{2.4}$$

The first term, R, is the loading rate, which is simply the rate at which atoms are captured from the background gas into the trap. The second term, γN, is the loss rate due to collisions between a trapped atom and a thermal atom from the background gas. The last term is also a loss rate, but one due to collisions between two trapped atoms. We would like to understand how the parameters here - R, γ and β - relate to the optimization targets discussed in sec. 2.1. While γ depends mainly on characteristics of the background gas, R is highly dependent on the forces acting on the atoms, and will provide the main connection between our performance measure and our optimization targets.

2.2.1 The Reif model

In more precise terms, R is the rate at which atoms that can be captured enter the trap volume. An atom can be captured when it has velocity v lesser than some capture velocity v_c. Let V be the trap volume. Our problem may then be formulated as such: what is the flux of atoms with $v \leq v_c$ through the surface of V? The answer is given as what has been called the Reif model.[9]

Consider an area element dA on the surface of V, and let our coordinate frame be defined in such a way that the z axis is parallel to the external normal vector dA. Let $f(\mathbf{v})d^3v$ be the fraction of atoms with velocity between $\mathbf{v}$ and $\mathbf{v} + dv$, and n the atomic density. The atomic flux per unit area over the surface of V for atomic velocities between $\mathbf{v}$ and $\mathbf{v} + d\mathbf{v}$ is then $\phi\, d^3v = nf(\mathbf{v})v\cos\theta d^3v$. Then the flux per area of atoms that can be captured is found by integrating $\phi\, d^3v$ for $v \leq v_c$ and $v_z > 0$, which already excludes atoms with velocity pointing away from the trap, i.e., ones that do not cross its surface. For a gas in equilibrium, we have $f(\mathbf{v}) = f(v)$. Then, by using the Maxwell-Boltzmann distribution for $f(v)$ we arrive at an exponential solution which expanded to second order on v_c^2 gives

$$\Phi = \frac{n}{4\pi^{1/2}} \frac{v_c^4}{v_{\text{mp}}^3} + \mathcal{O}(v_{mp} v_c^6), \tag{2.5}$$

where $v_{\text{mp}} = \sqrt{2k_BT/m}$ is the most probable velocity of the distribution. The loading rate R may then be obtained as $A\Phi$, the surface area of the trap, but it is usually expressed in terms of the trap volume V. So for low capture velocities, we have a good approximation[4]

$$R = \frac{nV^{2/3}}{2}\frac{v_c^4}{v_{mp}^3}, \tag{2.6}$$

2.2.2 The loss rate

We now turn our attention to the two loss mechanisms in the trap. Losses due to collisions with the background gas are proportional to N, as given by the collision rate γ. These collisions very easily eject atoms from the trap because these traps usually have depths of less than 1 K while the walls of the chamber are at room temperature.[10] The rate γ can be expressed in terms of the background gas temperature T_{bg}, density n_{bg} and the cross-section σ for collisions that eject atoms from the trap:[4]

$$\gamma = n_{\text{bg}}\sigma\sqrt{\frac{3k_BT_{\text{bg}}}{m}}, \tag{2.7}$$

where m is the mass of the atoms involved. Naturally, the warmer and denser the background gas, the greater the contribution of these collisions to the total loss rate. The cross-section σ is determined experimentally, but because the atom only leaves the trap if it gains a velocity higher than the capture velocity, γ should be smaller for greater v_c. Theoretical estimates have placed this dependence as $\sigma \propto v_c^{-2/3}$, which is very slight for typical values of v_c, on the order of 10 m/s.[10]

Losses due to collisions in between the trapped atoms present a somewhat more complicated problem, as they are, in fact, light-assisted collisions. Because they involve pairs of cold atoms, collision times become comparable to the lifetime of the excited state, and as a consequence absorption and emission may play a defining role in them. The dependence of this term on the *square* density in eq. (2.4) means that it only becomes significant for high densities; β itself depends on light intensity and detuning, but not density.[6] Intratrap collisional losses are therefore more significant in the low-atom "ideal gas mode", where density is proportional to N and shows a Gaussian profile with high density maximum. In contrast, for high populations, double scattering and beam attenuation become significant, resulting in a net repulsive force which balances with the trapping force, making density constant and uniform across the cloud even as N varies.[5]

Under these conditions, we are reasonably justified in dropping the contribution of intratrap collisions to the loss rate. The first point to consider is that a modern magneto-optical trap operates on very high trap populations in comparison to the works that originally treated the problem of density; for the trap we will be optimizing, loading occurs with atom numbers on the order of 10^8, for which density is uniform across the cloud.[11] In this situation, eq. (2.4) becomes $\dot{N} = R - \gamma N - \beta n^2 V$, where

V is the volume of the atomic cloud and n its density. As a second point, for such high atom numbers we are working in the regime where the contribution of the β term is expected to be very small; furthermore, when working with a vapor-cell MOT, n_{bg} is particularly high. Given those conditions, we conclude that γ dominates over β, so

$$\frac{dN}{dt} = R - \gamma N \tag{2.8}$$

This is the rate equation most frequently found in studies of the loading dynamics of a MOT.[4] We have nonetheless sought to justify why the experimental conditions are such that it may be used on the basis of the original publications where this was motivated.

2.2.3 The loading curve model

We may now actually derive two models from (2.8). For small N, we may take $R - \gamma N \approx R$. This gives us a simple linear evolution of the trap population at a constant rate R, for an initial loading period when there are still too few trapped atoms for collisions to take place and play a significant role. Beyond that, for the curve as a whole we have

$$N(t) = \frac{R}{\gamma}(1 - e^{-\gamma t}), \tag{2.9}$$

These models leave us with two physical parameters measured from loading curves to work with: the loading rate R and the background loss coefficient γ. Measuring these quantities is a simple matter of fitting eq. (2.9) to the data, as done in Fig. 2. R may be determined beforehand by fitting a linear model to initial loading times. Our model also tells us that for large t, we have a steady state population $N_f = R/\gamma$. This is what allows us to estimate the final trap population without actually loading the MOT up to the steady state, as in Fig. 2. As long as we measure the loading process until the loss contribution is clearly visible, we are able to fit the curve to eq. (2.9) and obtain R, γ and N_f.

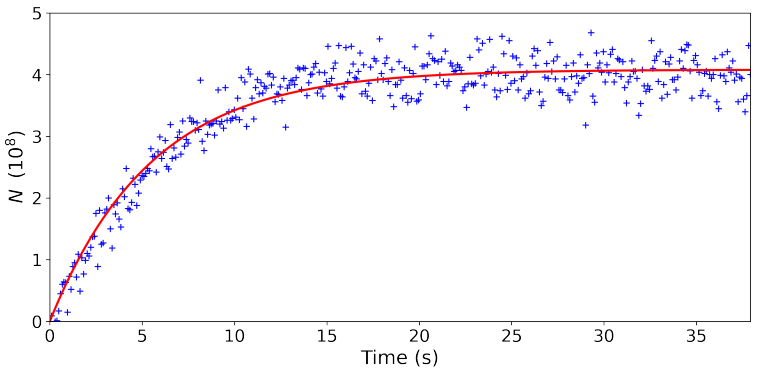

Figure 2: The solid line (red) shows the model in eq. (2.9) fitted to typical experimental data points. Source: By the author.

3 Machine learning optimization

3.1 Implementation

The goal of an optimization procedure is most simply put as the minimization of some quantity. In principle, any process may be optimized as long as its quality can be measured as a number. More precisely, optimization is the problem of determining the global minimum of a function $f(\mathbf{X})$ over an M-dimensional parameter space, $\mathbf{X}$ a vector of that space. This *cost function* f must then be properly defined so that its minimization implies in the desired optimization.

As established in the previous section, we measure the performance of a particular loading curve on the basis of its R and γ parameters, which may be obtained from a theoretical fit. As a result, we now have a means of estimating the final trap population N_f without fully loading the trap. Accordingly, we define our cost function to be

$$f(\mathbf{X}) = \frac{1}{N_f} = \frac{\gamma}{R} \tag{3.1}$$

With the vector $\mathbf{X}$ being the 4-dimensional vector in the space of the target parameters discussed in sec. 2.1. We have already shown in sec. 2.2.1 that R is strongly related to the capture velocity ($R \propto v_c^4$), which is directly determined by the scattering force.[7] On the other hand,γ depends on the capture velocity slightly, but because collisions with thermal atoms are likely to eject atoms in any case, is mostly determined by the background gas. We then expect optimization of N_f to generally be followed by an increase in R, viewed as an increased initial capture rate, but not by a considerable decrease in γ, which would be visible for a long curve as a delayed onset of saturation.

The learner has some limited knowledge of the problem at the start of the process. We supply it with an initial set of parameters $\mathbf{X}_0$, as well as bounds on the four parameters. These bounds are chosen not only to keep the learner within acceptable configurations for the experimental apparatus, but also to not let it stray too far into regions of the parameter space that produce no loading; some amount of capture must happen, or otherwise there will be no actual loading curve from which an accurate cost can be determined. The learner was configured to run for a set number of executions of the experiment. Then the working cycle is as follows: the learner selects a set of parameters and writes them to a configuration file, which is then read by the experimental control software; the MOT loading is performed for a predefined time interval (of the order of 10 s), and the loading curve is stored; then the cost is measured from the curve and returned to the learner; the learner is then updated with this information and selects another $\mathbf{X}$.

For implementing the learner itself we have used the `M-LOOP` toolkit introduced in sec. 1, an open

source tool based on the Python library `scikit-learn`, which has been successfully used for the optimization of different stages of BEC production sequences[12–14] It may be run from the command line or implemented as a Python API. Here we have opted for the latter in order to maintain flexibility.

3.2 Optimization methods

Before describing the specific methods implemented, we should go over the manner in which they are operated. Because the learner here chooses the parameter sets to be tested at each run, it can behave in two distinct ways. It might act as the perfect ”optimizer”, at each run selecting only the optimal parameters and using the returned cost function to update this model; this guarantees convergence to a cost minimum, but also means that it would not be able to avoid local minima. Alternatively, it might act as an ”explorer”, by undertaking a wider exploration of the parameter space in order to collect a higher variety of data which will allow it to build a more accurate picture of how parameters relate to cost; this would allow it to properly identify the global minimum given enough time, but it would require a larger number of runs in order to actually explore the whole parameter space and lack refinement in determining the minimum. Since we aim here to build an efficient and precise learner, the ”optimizer” approach is more adequate, but some amount of exploration should still be considered in order to avoid local minima.

`M-LOOP` tackles this problem by simultaneous use of two learners. One is a pure ”optimizer”, as described. The other is the ”explorer” responsible for generating an initial training set for the optimizer and providing the optimizer with unbiased data at regular intervals. This approach seeks to reach a good compromise between accuracy and sufficient exploration of the parameter space.

3.2.1 Differential evolution

Differential evolution (DE) is an evolutionary algorithm, and so its basic principle is in testing a set of individuals (parameter sets) at each generation, selecting the fittest (lower cost) and combining and/or mutating them in order to produce the next generation. In DE, a new individual $\mathbf{X}'$ is produced by a process called mutation: adding the weighted difference between two individuals of the previous generation to a third individual $\mathbf{X}$. Then crossover is applied in order to increase diversity: a trial individual $\mathbf{Y}$ is built by combining randomly selected parameters from $\mathbf{X}$ and $\mathbf{X}'$. $\mathbf{Y}$ replaces $\mathbf{X}'$ in the new generation only if it incurs a cost improvement. Because many individuals spread out across the parameter space are tested and then mixed before applying this ”greedy” criterion, individuals closer to the global minimum might help others escape local minima.[14] For this reason, differential evolution acts as the ”explorer” and generates initial training sets for the other methods.

3.2.2 *Gaussian process regression*

In general terms, a Gaussian process (GP) is a probability distribution over functions that describe a given dataset. In GP regression, the cost is modeled by a GP with constant mean function and covariance, defined by a *correlation function*. As implemented by `M-LOOP`, this is a Gaussian correlation function,[12] that depends on $H = (h_1, \ldots, h_M)$, the set of correlation lengths for each parameter. As K defines the mean and covariance, and is dependent on H, the correlation lengths act as hyperparameters of the model, which must be fitted online. We can consider the likelihood of parameters H given observations $\mathcal{O}$, and pick the most likely H. For small datasets, however, there might be multiple H with comparable likelihoods close to the maximum. We define from these H a hypothesis set $\mathcal{H} = (H_1, \ldots, H_P)$, with a corresponding likelihood set. Given a hypothesis set, a weighted average is performed over it in order to determine the final estimates of the mean function $M_{\hat{\mathcal{C}}}(\mathbf{X}|\mathcal{O}, \mathcal{H})$ and variance. For a set $\mathcal{O}$ of observations and a set $\mathcal{H}$ of correlation length hypotheses, the optimal $\mathbf{X}$ for $M_{\hat{\mathcal{C}}}(\mathbf{X}|\mathcal{O}, \mathcal{H})$ are selected for a new run, returning a cost $C(\mathbf{X})$. Then the pair $\mathbf{X}, C(\mathbf{X})$ is added to $\mathcal{O}$ and updates $\hat{\mathcal{C}}$, the learner's model of $\mathcal{C}(\mathbf{X})$, before selecting a new set of parameters.

3.2.3 *Artificial neural networks*

Being constructed as a network of nodes called neurons, an artificial neural network (ANN) maps a set of inputs $\mathbf{X}$ to an output Y, in a relation determined by the set of weights and connections in that network. `M-LOOP` implements a stochastic artificial neural network (SANN) comprised of three ANNs plus a DE learner that introduces the stochastic element to the overall ANN. After DE generates the initial training set, each ANN and the DE learner return a set of optimum parameters, which are tested in turn. The four resulting parameter-cost pairs are then used to update all three ANNs. Having three ANNs in addition to DE helps ensure that the parameter space is properly explored, once again decreasing the chance that optimization will get caught in a local minimum.[13]

4 Experimental setup

The optimization procedure previously discussed has been applied to the cooling and trapping of ^{87}Rb atoms in a vapor-cell magneto-optical trap. The relevant transitions are in the D2 line at 780 nm.

The cooling beams are generated by a *Toptica TApro* laser and the repump beam is generated by a *Toptica DLX100* laser. The cooling beams are amplified by a homemade MOPA (Master Oscillator Power Amplifier). Both lasers are locked at about 780 nm wavelength within 1 MHz tolerance. The precise wavelengths of each laser are stabilized via saturated absorption spectroscopy. The typical powers used for the cooling beams are 270 mW trapping and 30 mW repumping.

The control system uses two *National Instruments* data acquisition cards, the *PCI 6359* and the *PCI 6733*. The experimental control sequence is written in Python and executed through a *LabVIEW* virtual instrument (VI). Our optimizer automatically writes the selected parameter set to the text file which sets the parameter values before each run.

The quadrupole magnetic field, as has already been mentioned, is generated by a pair of coils in the anti-Helmholtz configuration. The compensation coils are arranged in three pairs, each one in the Helmholtz configuration along a different axis. The magnetic field gradient is usually set to about 10 G/cm.

For the measurement of the loading curve, fluorescence from the trapped atomic cloud generates a voltage drop in a photodetector, which is recorded over time. The voltage V_f is linear with the optical power received, with a responsivity $\mathcal{R}$. We can then establish a linear relation between the number of trapped atoms N and V_f. If the cooling transition wavelength is λ and the scattering rate is the intensity-dependent Γ_s as given by (2.1), the total power radiated by a single atom is

$$P = \frac{h\Gamma_s c}{\lambda} \tag{4.1}$$

Taking the fluorescence to be isotropic, the radiated power per solid angle is simply $p = P/(4\pi)$. Then the power collected by the photodetector from N atoms is $P_f = p\Omega_f N$, where Ω_N is the solid angle covered by the photodetector. The relation between N and I_f is then

$$N = \frac{1}{\Gamma}\frac{V_f/\mathcal{R}}{hc/\lambda}\frac{4\pi}{\Omega_f} = \kappa V_f \tag{4.2}$$

For the photodetector utilized here, we have determined $\kappa = 1.1 \times 10^8\, V^{-1}$ after calibration.

5 Results

We have applied the implementation described in sec. 3.1 to the loading of the magneto-optical trap described in sec. 4. In every optimization run, differential evolution (DE) was used to generate an initial training set of 11 loading runs. We have tested optimization for Gaussian process regression (GP), an artificial neural network (ANN) and DE. Optimization was performed over 150 runs for a loading time of 15 s each for GP, ANN and DE. The magnetic field is automatically turned off at the end of each run in order to completely empty the trap, and then turned back on again before the next run. Under these conditions, the optimization took about 45 minutes.

For the starting set of parameters, we provided the learner with unoptimized parameters that were enough to produce a typical loading curve. These initial settings produced a cloud of 2.8×10^8 atoms, which is about three times smaller than the best human optimization. These initial settings reflect what might be the daily use of ML: the optimizer need only be set to an approximately good starting configuration, and then it takes care of fine-tuning parameters for the specific conditions of the experiment.

Excluding from the count the 11 initial training runs, in the optimization above the ANN learner converged to cost 0.018 in 70 runs, while the GP learner converged to cost 0.041 in 14 runs and the DE learner achieved cost 0.096 in 114 runs. The slow convergence of DE is what motivates implementing it solely as an "explorer" to support the more robust methods. Although the cost reached by GP is not so much lower than that reached by DE, the GP took 100 fewer runs. However, the GP was not as effective as the ANN, which reached the lowest cost by a fair margin, although at a slower pace.

To assess the effectiveness of our proposed loading curve method, and more clearly compare the results of the ANN and GP learners, we loaded the MOT with their best optimized parameters for 35 s, and then compared the results with the best human optimization, as shown in Fig. 3. Human optimization reached a trap population of 7.84×10^8 atoms, while the GP method reached 8.02×10^8 and the ANN method 8.25×10^8, resulting in respective gains of 1.8×10^7 and 4.1×10^7 in relation to the human optimization. These are considerable amounts, because further cooling stages in BEC production sequences involve allowing the escape of the most energetic atoms from the cloud, and a good BEC generally has atom numbers on the order of $1 - 2 \times 10^5$. Atom numbers on the order of 10^7 are actually the minimum trap population necessary for the BEC production in the lab, being enough even when under conditions that lead to the loss of 3/4 of the initially trapped atoms.

Fig. 4 shows the cost as a function of run number for every run of the experiment for each learner. Shown alongside the experimental points is a solid line representing the evolution of the minimum

cost from all parameter sets tested up to each run, i.e., at each run number the solid line indicates the minimum of the set of all costs obtained up to that particular run. The solid line highlights the convergence of the cost towards a minimum in every case, at clearly distinct rates for each learner.

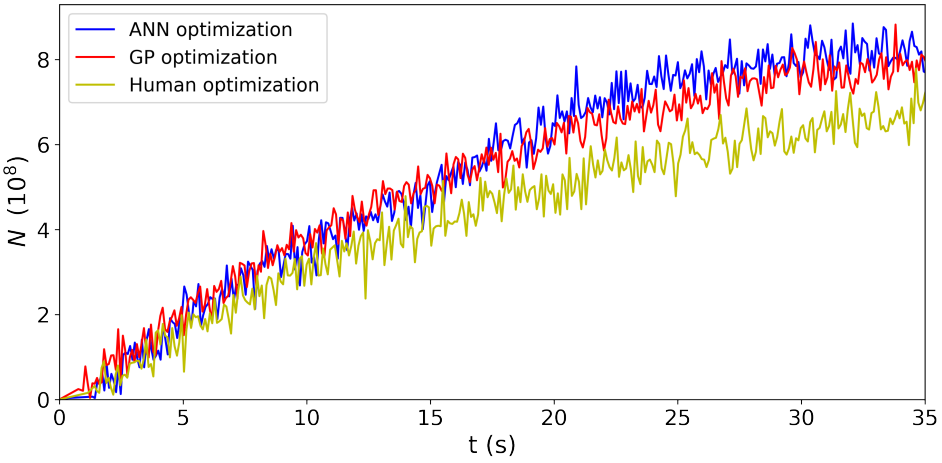

Figure 3: Comparison between the best optimized curves for ANN, GP and a typical human optimization.
Source: By the author.

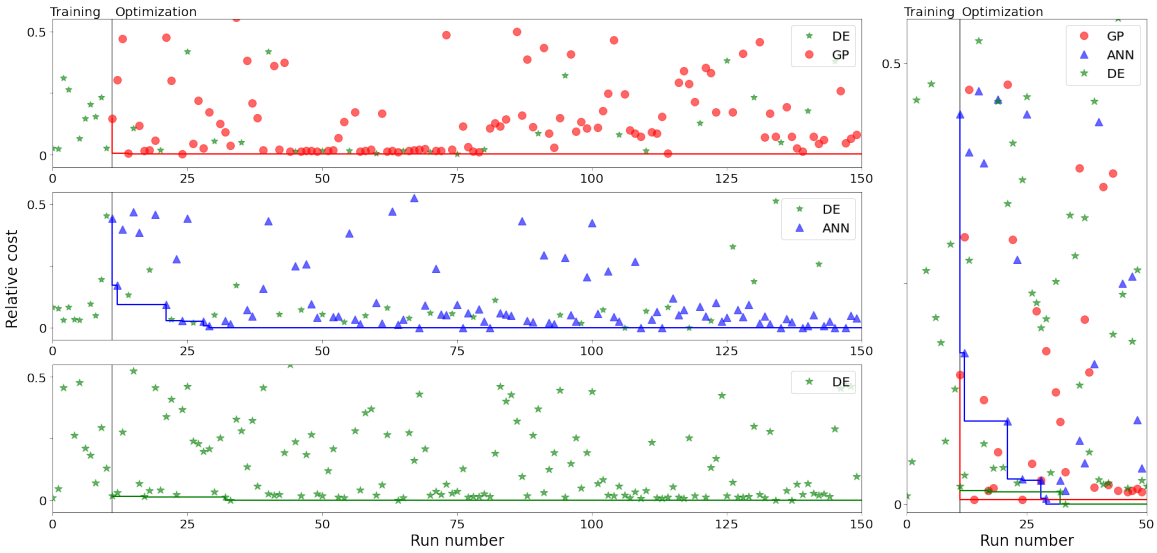

Figure 4: Cost per run, by learner, normalized by the overall minimum and maximum costs obtained. The first 11 points form the training set generated by DE. The solid line shows the evolution of the minimum cost found by each learner. On the left we have the GP (top), ANN (mid) and DE (bottom) learners, while the inset on the right shows the first 50 runs for all algorithms.
Source: By the author.

6 Conclusions

Our proposal for the performance measure was successfully shown to be able of guiding machine learning algorithms towards the optimization of a magneto-optical trap exceeding the best human optimization. We have found a gain of approximately 4.1×10^7 atoms in relation to the best previous optimization of the experiment, nearly achieving a total of 10^9 atoms in the trap for loading times of 35 s, which are the usual loading times used in the laboratory.

We have also been able to compare different optimization methods. As expected, DE was the least efficient, reaching 150 runs to achieve the best cost, which is still twice as many as the GP and more than three times than the ANN. We have observed that GP converges to a low cost quickly, while the ANN, although slower, eventually reaches a better cost than GP, in agreement with the literature.[14] One possible reason for this is that ANNs are generally better able to deal with large data sets than GP, and so, as more runs are accumulated, the GP learner becomes progressively slower in comparison to ANN. The implementation of the ANN learner as a SANN which simultaneously uses three ANNs to explore the parameter space might also be more likely to find a global minimum in the long run, though both GP and ANN receive unbiased, "exploratory" data from DE. Overall, the relative convergence rate and optimization quality observed between ANN, GP and DE seems to be in agreement with previous work based on the same algorithms for other stages of BEC production.[14]

The high quality of the optimization results comes together with a relatively short optimization time. The best optimization was achieved by the ANN in 45 minutes of optimizing runs, but the actual best parameters were found by the learner in under 100 loading runs, after about 24 minutes of execution. This makes this implementation a practical tool for everyday use, as it is not only able to provide a better optimized MOT, but it can also be run often enough to keep the experiment close to peak performance every day. The short optimizations times also open up the possibility of reducing the number of runs while increasing the loading time of individual runs, allowing for an improvement in optimization precision. Although, in our runs, the GP was not able to provide the same level of optimization as the ANN, this extra balance comparison between the loading and optimization times might help to determine the best option.

References

1 ASHKIN, A.; GORDON, J. P. Stability of radiation-pressure particle traps: an optical earnshaw theorem. *Optics Letters*, v. 8, n. 10, p. 511, Oct. 1983.

2 PRITCHARD, D. E. *et al.* Light traps using spontaneous forces. *Physical Review Letters*, Woodbury, v. 57, n. 3, p. 310–313, July 1986.

3 RAAB, E. L. *et al.* Trapping of neutral sodium atoms with radiation pressure. *Physical Review Letters*, Woodbury, v. 59, n. 23, p. 2631–2634, Dec. 1987.

4 MONROE, C.; SWANN, W.; ROBINSON, H.; WIEMAN, C. Very cold trapped atoms in a vapor cell. *Physical Review Letters*, Woodbury, v. 65, n. 13, p. 1571–1574, Sept. 1990.

5 SESKO, D. W.; WALKER, T. G.; WIEMAN, C. E. Behavior of neutral atoms in a spontaneous force trap. *Journal of the Optical Society of America B*, v. 8, n. 5, p. 946, May 1991.

6 TELLES, G. D. *Estudo de colisões atômicas ultrafrias: mecanismos de perda e espectroscopia de fotoassociação.* 2002. p. 45. Tese (Doutorado em Ciências) - Instituto de Física de São Carlos, Universidade de São Paulo, 2002.

7 LINDQUIST, K.; STEPHENS, M.; WIEMAN, C. Experimental and theoretical study of the vapor-cell zeeman optical trap. *Physical Review A*, Woodbury, v. 46, n. 7, p. 4082–4090, Oct. 1992.

8 FOOT, C. J. *Atomic physics*. Oxford: Oxford University Press, 2005.

9 REIF, F. *Fundamentals of statistical and thermal physics*. New York: McGraw-Hill, 1965.

10 STEANE, A. M.; CHOWDHURY, M.; FOOT, C. J. Radiation force in the magneto-optical trap. *Journal of the Optical Society of America B*, v. 9, n. 12, p. 2142, Dec. 1992.

11 BARTALINI, S. *et al.* Full characterization of the loading of a magneto–optical trap from an alkali metal dispenser. *The European Physical Journal D*, v. 36, n. 1, p. 101–104, June 2005.

12 WIGLEY, P. B. *et al.* Fast machine-learning online optimization of ultra-cold-atom experiments. *Scientific Reports*, v. 6, n. 1, May 2016. DOI: 10.1038/srep25890.

13 TRANTER, A. D. *et al.* Multiparameter optimisation of a magneto-optical trap using deep learning. *Nature Communications*, v. 9, n. 1, p. 4360, Oct. 2018.

14 BARKER, A. J. *et al.* Applying machine learning optimization methods to the production of a quantum gas. *Machine Learning:* science and technology, v. 1, n. 1, p. 015007, Feb. 2020.