RESEARCH ARTICLE | SEPTEMBER 10 2025

Efficient production of sodium Bose–Einstein condensates in a hybrid trap

Y. Geng ; S. Mukherjee ; S. Banik ; M. Gutierrez Galan ; M. Anderson ; H. Sosa-Martinez ;
S. Eckel ; I. B. Spielman ; G. K. Campbell

Rev. Sci. Instrum. 96, 093202 (2025)
https://doi.org/10.1063/5.0279068

Review of Scientific Instruments | ARTICLE | pubs.aip.org/aip/rsi

Efficient production of sodium Bose–Einstein condensates in a hybrid trap

Cite as: Rev. Sci. Instrum. 96, 093202 (2025); doi: 10.1063/5.0279068
Submitted: 5 May 2025 • Accepted: 8 August 2025 •
Published Online: 10 September 2025

View Online Export Citation CrossMark

Y. Geng,[1] S. Mukherjee,[1] S. Banik,[2] M. Gutierrez Galan,[3] M. Anderson,[1] H. Sosa-Martinez,[4] S. Eckel,[5] I. B. Spielman,[1,5,a)] and G. K. Campbell[1,5,b)]

AFFILIATIONS

[1] Joint Quantum Institute, University of Maryland and National Institute of Standards and Technology, College Park, Maryland 20742, USA
[2] Lightmatter, Inc., Mountain View, California 94040, USA
[3] IonQ, College Park, Maryland 20742, USA
[4] Capital One, McLean, Virginia 22102, USA
[5] National Institute of Standards and Technology, Gaithersburg, Maryland 20899, USA

[a)] Author to whom correspondence should be addressed: ian.spielman@nist.gov
[b)] Electronic mail: gretchen.campbell@nist.gov

10 September 2025 13:05:21

ABSTRACT

We describe an apparatus that efficiently produces ^{23}Na Bose–Einstein condensates (BECs) in a hybrid trap that combines a quadrupole magnetic field with a far-detuned optical dipole trap. Using a Bayesian optimization framework, we systematically optimize all BEC production parameters in modest-sized batches of highly correlated parameters. Furthermore, we introduce a Lagrange multiplier-based technique to optimize the duration of different evaporation stages constrained to have a fixed total duration; this enables the progressive creation of increasingly rapid experimental sequences that still generate high-quality BECs. Taken together, our techniques constitute a general approach for refining and accelerating sequence-based experimental protocols.

Published by AIP Publishing. https://doi.org/10.1063/5.0279068

I. INTRODUCTION

Since their first experimental realization in 1995,[1,2] atomic Bose–Einstein condensates (BECs) have become a workhorse system for studying a diverse range quantum many-body phenomena,[3–5] including open quantum systems,[6] quantum gases in low dimensions,[7–11] and cosmological as well as gravitational physics.[12–20] This breadth of applications has been accompanied by generations of increasingly sophisticated apparatuses, each incorporating new techniques for cooling, control, and measurement. Quantum gas experiments nearly universally operate in a cyclic manner, where each repetition creates a quantum gas that is then manipulated, measured, and, in the process, destroyed. Cutting-edge experiments require very large datasets—making short cycle times highly desirable—and demand a high degree of stability with minimal drift in experimental conditions.

The overall cycle time is often dominated by the preparation process, which can take tens of seconds. Once the quantum gas is created, the subsequent physics experiment can take anywhere from a few microseconds to many tens of seconds. As a result, the preparation time is often the bottleneck for employing quantum degenerate gases in both applied and fundamental applications. An increased data rate enables more effective exploratory studies, yields larger datasets for reduced statistical noise, lowers Dick sampling noise[21] for precision measurements, and reduces sensitivity to systematic drifts. Thus, quantum gas production techniques that reduce production time and increase stability are essential for further applications of this platform. In this paper, we begin by briefly describing our experimental apparatus for generating ^{23}Na BECs; we then detail several strategies for optimally creating BECs while simultaneously minimizing the production time.

Quantum degenerate gases are almost always produced in a two-stage process consisting of laser cooling, followed by evaporative cooling. All such experiments include a magneto-optical trap (MOT) as part of the laser cooling process to prepare an initial sample of cold atoms. In some cases, such as ours, this MOT collects atoms from a cold atomic beam (e.g., from a Zeeman slower, a 2D-MOT, or a buffer gas source), while in others, atoms are captured directly from the low velocity tail of a dilute room-temperature vapor. In our experiment, the MOT temperature is fairly high (~290 μK), so we include a sub-Doppler cooling stage that reduces the temperature to about ≈40 μK before transferring the atoms to a conservative potential (here, a magnetic trap) for the evaporative cooling stage. Sub-Doppler cooling in ^{23}Na leaves the sample well above the typical ≈5 μK BEC transition temperature, thereby requiring efficient evaporative cooling. Evaporative cooling selectively removes the most energetic atoms while elastic collisions rethermalize the remainder at a reduced temperature. Because collision rates in these dilute gases are relatively low, evaporative cooling typically accounts for the majority of the total production time. A primary focus of this paper is, therefore, to optimize evaporation performance subject with a constrained total evaporation time.

Given the clear delineation between initial laser cooling and the subsequent evaporation, we quantify the evaporation stage's performance by introducing the production efficiency $N_{\rm BEC}/N_{\rm MOT}$. This ratio reflects how effectively a cloud of $N_{\rm MOT}$ laser-cooled atoms is converted into a fully condensed BEC of $N_{\rm BEC}$ atoms. Figure 1 plots production efficiency as a function of the evaporation time $t_{\rm evap}$ (rather than total cycle time) for various ^{23}Na BEC apparatuses, thereby isolating the effect of evaporation from any complexities in MOT loading. Because the MOT loading time strongly depends on variables such as the atomic source, laser power, and beam geometry, we treat it as largely independent of evaporation optimization. Nonetheless, in the outlook we briefly note how coupling these two stages could provide further performance gains.

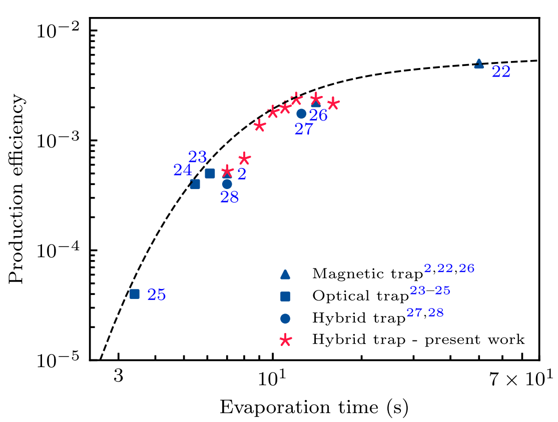

FIG. 1. Comparison of different sodium BEC machines' production efficiency and evaporation time. Different symbols identify the underlying trapping employed during evaporation: magnetic traps (triangles),[2,22,23] all-optical traps (squares),[24–26] and hybrid magnetic-optical traps (circles),[27,28] with our system represented by red). In addition, the dashed curve serves as a guide to the eye, indicating the nominal performance limit of existing systems in the efficiency-$t_{\rm evap}$ plane.

Harmonically confining magnetic traps generated by large, out of vacuum, coils feature very large trap volumes but with correspondingly gentle confinement.[22,29] As a result, they not only yield the largest $N_{\rm BEC}$ but also require large $t_{\rm evap}$ due to their low atom densities and corresponding low collision rates. These low densities decrease 2- and 3-body loss processes, contributing to their high efficiencies.

In contrast, purely optical traps[24–26,30] achieve very strong confinement; while this enables rapid evaporation, the resulting BECs are small due to limited trap volume, and enhanced three-body losses. Similarly in ^{87}Rb, magnetic chip traps[31,32] offer strong confinement in a compact geometry and generate small BECs with cycle times as short 1 s.

Hybrid traps, initially developed for ^{87}Rb, merge the strengths of magnetic and optical dipole traps (ODTs): first capturing atoms in a large-volume magnetic trap and subsequently transferring them to a tightly confining optical trap.[27,33–35] In Fig. 1, they occupy an intermediate performance regime, achieving large $N_{\rm BEC}$ together with modest $t_{\rm evap}$. Notably, Ref. 2 employed an optically plugged quadrupole magnetic trap that achieved similar performance, but was rapidly abandoned owing to alignment instability. In these cases, the same higher density that reduces the evaporation time also accelerates 2- and 3-body loss processes.

For all of these BEC production approaches above, BECs are produced through a sequence of experimental stages each with a set duration, during each of which parameters—quantifying the power and frequency of optical, microwave, and RF electromagnetic fields, the strength and direction of (near) DC magnetic fields, to name just a few—follow preselected ramps. Efficient BEC production, therefore, requires not only a carefully chosen hardware configuration but also careful optimization of the stage durations and parameter values. After manually establishing a preliminary parameter set, we employ a Bayesian optimization framework `M-LOOP` introduced in Ref. 36 that is known to be performant for cold atom experiments.[36–39] In isolation, Bayesian optimization suffers from the "curse of dimensionality," leading to slow convergence (or none at all) when more than a handful of parameters are involved. We address this by measuring the parameters' covariance matrix, grouping strongly correlated parameters, and then employing Bayesian optimization to these groups sequentially.

This approach is effective in generating the globally optimal parameter set; however, it is unable to support algebraic constraints, such as optimizing the duration of each experimental stage subject to the constraint of fixed overall sequence duration. We, therefore, developed a Lagrange multiplier scheme to enforce such constraints, enabling us to optimize the range of fixed-duration evaporation sequences whose performance is shown in Fig. 1. Combining these approaches, we produced large BECs of $4.2(1) \times 10^6$ atoms, starting with $1.8(1) \times 10^9$ atoms in the MOT, after 12 s of evaporation, and still yielded $9.2(1) \times 10^5$ atom BECs using just 7 seconds of evaporation.

The remainder of the manuscript is organized as follows: in Sec. II, we provide an overview of our experimental design at the hardware level; then in Sec. III, we outline our BEC production sequence; next, Sec. IV details our optimization protocol; and finally, we conclude in Sec. V.

II. SYSTEM DESCRIPTION

Here, we provide a high-level description of our apparatus, including descriptions of the mechanical and optical setups. Further details can be found in the Ph.D. thesis of Banik[40] and Gutierrez Galan.[41]

A. Vacuum system

The vacuum system (Fig. 2) comprises a "high" pressure 2D MOT source, connected via a low conductance link to a low-pressure UHV science chamber (giving a vacuum-limited lifetime of magnetically trapped atoms ≈25 s).

1. 2D MOT chamber

Our 2D MOT (Fig. 3), inspired by Ref. 42, is built around a ten-way ConFlat cross with four 2.75 in. and four 1.33 in. flanges in the cross's plane, plus individual 2.75 and 1.33 in. flanges on the axis normal to that plane (1 in. = 2.54 cm). The 2.75 in. out-of-plane flange connects to the pumping section, and the 1.33 in. flange connects to science chamber. Finally, a gate valve and a differential pumping tube with 3 mm diameter separate the atomic source region from the science chamber. A resistively heated sodium oven is mounted on the bottom in-plane 1.33 in. flange.

The 2D MOT's quadrupole field is produced by four stacks of nine $25 \times 10 \times 3\ \mathrm{mm}^3$ neodymium bar magnets [Eclipse N750-RB,[43] with magnetization $(8.8 \pm 0.1) \times 10^5\ \mathrm{A\ m}^{-1}$] mounted on the cross.[42,44] The stacks are divided into front and back sets spaced by 70 mm, and each set contains two stacks that are vertically separated by 98 mm. These result in a 2D quadrupole field with a 0.42 T/m gradient; this gradient changes by only ≈10 % over the longitudinal extent of the 2D MOT.

2. Science chamber

The science chamber (Kimball Physics MCF800-SphSq-G2F1E3C4A16) is an 8 in. stainless-steel vessel with recessed windows on the top and bottom. Four 4.5 in. and four 2.75 in. flanges reside on the horizontal plane and eight pairs of 1.33 in. flanges are placed 21° above and below the horizontal plane. The recessed

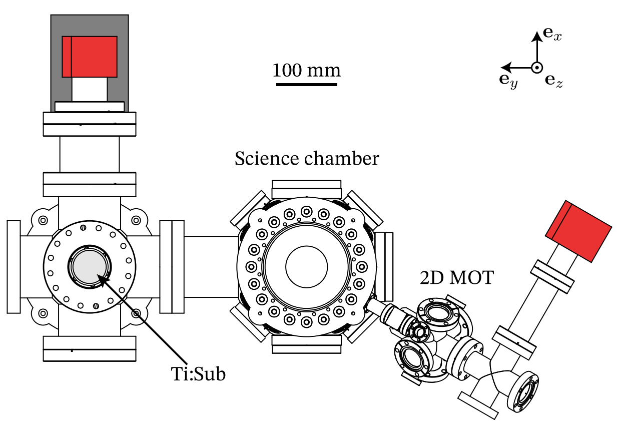

FIG. 2. Top view of the vacuum system. The red/gray colored regions denote vacuum pumps with a standard ionization pump in dark gray, combination ionization/getter pumps in red, and a Ti:Sub pump in light gray.

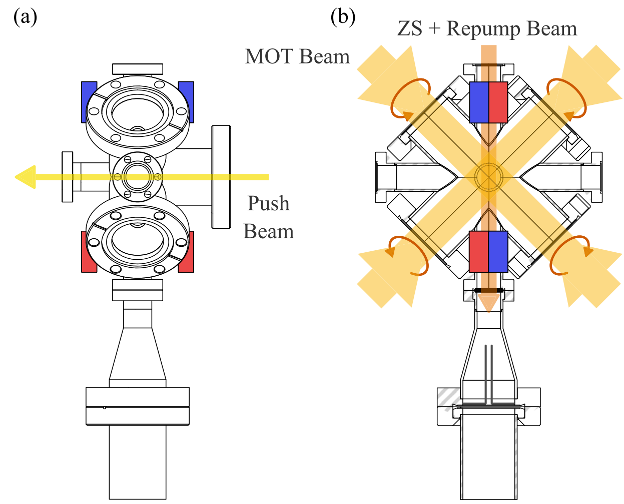

FIG. 3. 2D MOT. (a) Front view: laser beams are colored in yellow. The north and south poles of magnets are colored in red and blue, respectively. (b) Side-section view: the polarization is marked on two MOT beams. Zeeman slower beam is polarized perpendicular direction of the drawing.

windows' inner surfaces are setback by 20 mm from the chamber's center.

Magnetic field gradients for the 3D MOT and magnetic trap are generated by a pair of coils in an anti-Helmholtz configuration. Each coil is wound from Kapton insulated 3/16 in. (0.48 cm) hollow square copper tube and is mounted just outside a recessed window. At the peak current of 200 A, these coils provide a 2.2 T/m gradient along $\mathbf{e}_z$ and are water-cooled with a flow rate of $0.6\ \mathrm{Lmin}^{-1}$, supplied by a booster pump operating at 1.25 MPa (180 psi). Finally, three additional pairs of coils in a Helmholtz configuration provide bias magnetic fields along each Cartesian axis.

The 2D MOT chamber is actively pumped by a hybrid ion-getter pump (SAES D-100), and the science chamber is pumped by a combination of an ion pump (Gamma Vacuum 45S) and an hybrid ion-getter pump (SAES D-500). A titanium sublimation pump is attached to the science chamber but was only activated during the initial pump-down phase.

B. Optical sources

All laser light used for laser cooling, imaging, and repumping is derived from two lasers resonant with distinct transitions within ^{23}Na's 589 nm D_2 line. The first, a Toptica DL-RFA-SHG pro (1.1 W output power), addresses the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cooling transition. The second, a Toptica TA-SHG pro (0.9 W output power), couples to the $|F = 1\rangle \rightarrow |F' = 2\rangle$ repumping transition. The cooling laser is locked to a Doppler-free saturated absorption feature, while the repump laser is beat note-locked to the cooling laser with a 1.8 GHz frequency shift. All laser light is delivered to the vacuum system via single-mode optical fiber. The nominal laser parameters are summarized in Table I.

1. 2D MOT lasers

Our 2D MOT utilizes in-plane "cooling" beams, a repump beam, a Zeeman slower beam, and a transverse push beam. The two

TABLE I. Set of laser parameters during each laser cooling stage. Frequencies are reported with respect to that of the desired transition (i.e., $f_{2\to3}$ denotes the frequency of the $|F = 2\rangle \to |F' = 3\rangle$ transition). Intensities are reported with respect to the saturation intensity of the $|F = 2\rangle \to |F' = 3\rangle$ transition. The red detuning of the 3D MOT- and bright-repumps is a legacy of initial manual optimizations that found these values to be insensitive at the scale of a few linewidths.

Beam	Frequency	Intensity
MOT loading		
2D MOT cooling	$f_{2\to3} - 1.2\Gamma$	$1.9I_{sat}$
2D MOT repump	$f_{1\to2} - 1.5\Gamma$	$30.5I_{sat}$
Zeeman slower	$f_{2\to3} - 18.6\Gamma$	$23.1I_{sat}$
Push beam	$f_{2\to3} + 0.3\Gamma$	$0.05I_{sat}$
3D MOT cooling	$f_{2\to3} - 1.4\Gamma$	$1.7I_{sat}$
3D MOT repump	$f_{1\to2} - 1.1\Gamma$	$1.9I_{sat}$
Sub-Doppler cooling		
3D MOT cooling	$f_{2\to3} - 3.4\Gamma$	$0.3I_{sat}$
Bright repump	$f_{1\to2} - 1.1\Gamma$	$0.1I_{sat}$
Optical pumping		
3D MOT cooling	$f_{2\to3} - 0.9\Gamma$	$0.3I_{sat}$
Optical pumping	$f_{1\to1} + 0.26\Gamma$	$0.007I_{sat}$

retro-reflected cooling beams (≈6 mm waist) enter and exit the 2D MOT chamber via the four in-plane 2.75 in. viewports. The repump and the Zeeman slower beams are overlapped, enter the chamber through the top 1.33″ viewport, and pass through the 2D MOT before entering the Na oven. Finally, the push beam has a 1 mm waist and enters the chamber by the longitudinal axis 2.75 in. viewport (i.e., oriented along the longitudinal axis of the 2D MOT directed into the science chamber).

2. MOT lasers

A 2-in, 6-out fiber beam splitter (Evanescent Optics) then distributes light injected into a single fiber to the six outputs that source the 3D MOT/sub-Doppler cooling beams with waist ≈12 mm). Repumping during dark SPOT MOT operation is provided by a single beam (with waist ≈8 mm) with a 9 mm diameter dark spot imaged onto the MOT center. In addition, a single "bright" (i.e., no dark spot) repump illuminates the atoms during sub-Doppler cooling.

3. Optical pumping laser

The optical pumping beam is derived from the same laser system as the MOT repump light, frequency-shifted by an acousto-optical modulator to address the $|F = 1\rangle \to |F' = 2\rangle$ transition. It is sent along the quantization axis set by a uniform bias magnetic field ≈0.1 mT with a waist of ~4 mm and power 40 μW.

4. ODT laser

The 1064 nm ODT beam is produced by a ytterbium fiber laser (IPG YLR20-1064-LP-SF). After pre- and post-conditioning along with coupling through a polarization-maintaining photonic crystal fiber (NKT LMA-PM-10) with ≈70 % efficiency, 6.3 W of power is delivered to the vacuum system. The beam then enters through the 4.5 in. viewport on the y axis focused to a $1/e^2$ waist of 22 μm, forming an ODT with a maximum depth of 200 μK.

C. UV desorption source

We employ light-induced atomic desorption (LIAD) to inhibit Na buildup on the top 2.75 in. 2D MOT windows.[45,46] This light is sourced from a pair of LED lamps (Thorlabs M365LP1) and is merged into the MOT beams via dichroic beam splitters yielding 30 mW at each viewport.

III. SYSTEM OPERATION

In brief, our BEC production begins with a 2D-MOT that generates a cold atomic beam from a hot effusive ^{23}Na source. The cold beam then travels into the science chamber where ^{23}Na atoms are captured into a dark spontaneous-force optical trap (SPOT,[47] although duplicative we will conform to the standard language of "dark SPOT MOT"). Afterward, they are sub-Doppler cooled[48] and finally optically pumped into the magnetically trappable low field seeking state $|F = 1, m_F = -1\rangle$.

We abruptly turn on a quadrupole magnetic trap, where the atoms undergo rf-evaporation. Once the cloud becomes sufficiently cold and dense, it is transferred into a horizontally aligned single-beam ODT.[33] Forced evaporation by lowering the ODT power brings the sample to quantum degeneracy.

A. 2D MOT

The Na oven, operating at ≈240 °C, directs an up-going atomic beam into the 2D MOT. Owing to Na's low vapor pressure, this can rapidly coat the top 1.33 in. viewport, which we prevent by continuously heating the top viewport to 130 °C. Over longer timescales, Na also accumulates on the in-plane 2.75 in. 2D MOT cooling viewports. We prevent this accumulation via nightly exposure to the UV LIAD light.

The decreasing fringe field from the 2D MOT permanent magnets resembles the field profile of a traditional Zeeman slower as it rises from 0 mT to a maximum of 20 mT over a distance of about 50 mm. Motivated by this observation,[42] our cooling process begins by Zeeman—slowing the Na atomic beam using a red-detuned down-going slower beam. The slowed atomic beam is then captured and cooled into the 2D MOT before being directed toward the 3D MOT by the blue-detuned push beam. This configuration relies on a single repump beam—co-propagating with the slower beam and detuned by 9 MHz from the $|F = 1\rangle \to |F' = 2\rangle$ transition—for both the slowing and 2D cooling stages. We operate this beam at fairly high power to provide effective repumping for both, given the poor transmission of the viewport.

B. Dark SPOT MOT

After entering the science chamber, the cold atomic beam from the 2D MOT is captured by the dark SPOT MOT. During this stage, the quadrupole magnetic field has a z-axis gradient of 0.06 T/m. The three pairs of counterpropagating MOT beams are red-detuned from the $|F = 2\rangle \to |F' = 3\rangle$ cycling transition and the dark-SPOT repump beam are red-detuned from the $|F = 1\rangle \to |F' = 2\rangle$ transition. In the dark volume at the interaction of these beams, atoms accumulate in the ground state $|F = 1\rangle$ manifold, reducing the radiation pressure and increasing the maximum atomic density.[47]

All together, we observe an initial atom loading rate of 0.8×10^9 s^{-1} and the MOT number saturates at 2×10^9 after about 10 s. In practice, we load the MOT for 7 s, yielding 1.8×10^9 atoms at ≈290 μK.

C. Sub-Doppler cooling

After MOT loading, we perform sub-Doppler cooling by removing the magnetic field gradient; further red-detuning the MOT beams and reducing their intensity; and use the bright repump beam for uniform repumping. This is commonly known as polarization gradient cooling (PGC) in other literature. During this stage, stray magnetic fields are canceled using the bias coils. The atoms cool to ≈40 μK in just 2 ms and are fully depumped into the $|F = 1\rangle$ manifold.

D. Hybrid trap loading

Prior to magnetic trapping, the sub-Doppler cooled cloud nominally uniformly populates all three magnetic sublevels of the $|F = 1\rangle$ manifold; of these states, only the $|F = 1, m_F = -1\rangle$ is a magnetically trappable weak-field seeking state. We, therefore, optically pump into $|F = 1, m_F = -1\rangle$ in 1.2 ms using a low intensity (0.02 mW/cm^2) circularly polarized optical pumping beam driving the $|F = 1\rangle \to |F' = 1\rangle$, aligned to an ≈0.1 mT bias field along with depumping from the MOT beams.[22]

After optical pumping, we suddenly energize the quadrupole trap to a gradient of 0.73 T/m, capturing 60% of the atoms. The trapped ensemble is then compressed by increasing the gradient to 1.9 T/m in 240 ms, adiabatically increasing the temperature to ≈230 μK, at a nominally fixed phase space density of 10^{-5} [Fig. 4(d)] The ODT is then turned on to its peak power of ≈6 W; at this time, it contributes only a small perturbation (trap depth ≈200 μK) to the magnetic trap.

E. Evaporative cooling

After laser cooling, we use evaporative cooling to bring the atoms to quantum degeneracy. Figure 4 shows our optimized evaporation sequence (see Sec. IV) that leads to our largest BECs.

First, we perform rf-evaporation in a purely magnetic quadrupole trap, which uses an oscillating magnetic field. If the atoms' Zeeman shift equals the frequency of the rf magnetic field, their internal state is changed from $|F = 1, m_F = -1\rangle$ to $|F = 1, m_F = 0, +1\rangle$ and they are ejected from the trap. At the beginning of the rf-evaporation stage, our initial frequency of 48 MHz is quickly swept from 48 to 34 MHz in 0.05 s and then slowly reduced to 11 MHz in 4.8 s, resulting in a cloud of $\sim 50 \times 10^6$ atoms with temperature ≈80 μK.

Near the end of the rf evaporation process, we reduce the magnetic confinement (green region in Fig. 4), allowing the ODT (placed just below the quadruple center) to provide vertical support against gravity. Because the ODT contributes a dimple-like potential to the horizontal magnetic confinement, even adiabatic transfer leads to the increase in phase-space density,[49] as seen in Fig. 4. During the hybrid trap evaporation stage (blue region in Fig. 4), the ODT depth is ramped down from ≈200 μK to 17 μK following an exponential ramp and a linear ramp with combined duration of 4.2 s. About halfway through this stage, the phase-space density of the atoms crosses the critical phase density of 2.6, at which point a BEC first forms.[50] The sequence concludes with an almost-pure BEC, with $4.2(1) \times 10^6$ atoms confined in a hybrid trap with frequencies: $\omega_\rho = 2\pi \times 320$ Hz (radial confinement from ODT) and $\omega_x = 2\pi \times 22$ Hz (longitudinal confinement from combined magnetic and optical potentials).

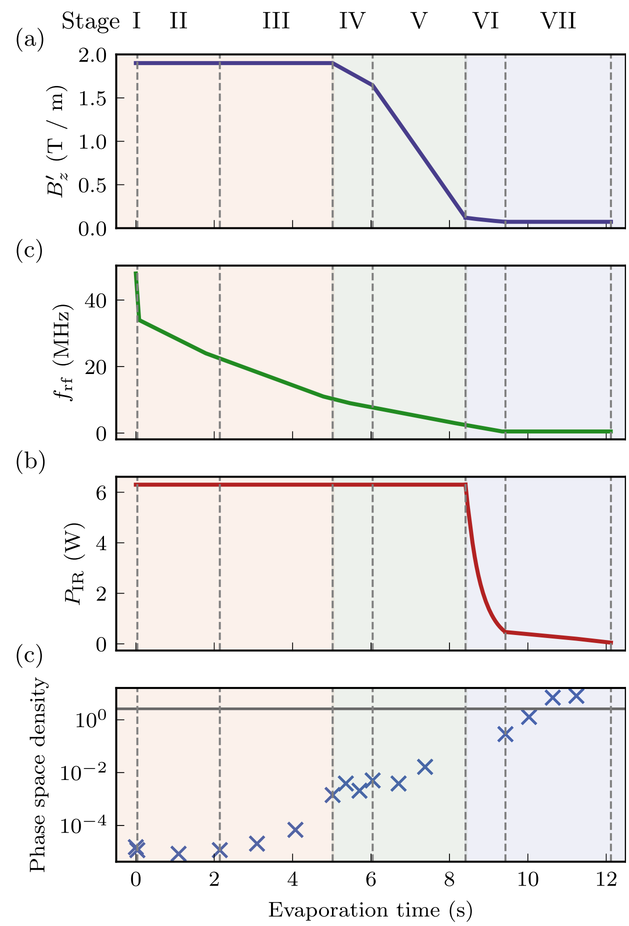

FIG. 4. Evaporation sequence: z-axis quadrupole field gradient B'_z, rf frequency $f_{\rm rf}$, and ODT beam power $P_{\rm IR}$ as a function of time into the evaporation sequence are shown in panels (a)–(c). The sequence is divided into ramps, whose boundaries are indicated by the gray dashed lines. The red, green, and blue shaded regions represent the rf evaporation stage, the decompression/transfer stage, and the hybrid trap evaporation stage, respectively. Panel (d) shows the phase space density calculated from the experimental data. The gray solid line marks the BEC transition's critical phase space density of ≈2.6.

IV. OPTIMIZATION

As seen in Sec. III, each separate experimental stage requires a range of control and timing parameters that are often correlated between stages. All of these interdependent parameters must be optimized for peak performance. Optimization can be abstracted as minimizing one or more "cost functions" that characterize the system's performance for a particular a set of control parameters.

A good selection of cost function requires it to sensitively reflect the system performance, while also being practically easy to acquire and robust to experiment noise. For us, we use the BEC fraction shortly after crossing the BEC phase transition (a point in the final ramp) as the cost function, as this value serves a sensitive agent of the phase space density. It can be reliably acquired by a single-shot measurement after a small time-of-flight, since the cloud is already cold but also has a reasonably small OD, in contrast to a pure BEC whose OD is difficult to determine. We note that the start and end points of the final ramp have to remain fixed for the measured BEC fraction to be comparable. To optimize the final ramp, we append an "anti-evaporation" ramp after the final ramp to adiabatically bring the system back to a fixed state and use the BEC fraction there as the cost function.

For experiments such as ours, brute force optimization quickly becomes intractable in the face of a large number of strongly correlated control parameters, an associated convoluted cost function landscape, measurement noise, and of course hardware-level technical limitations. It is virtually impossible to construct a top-to-bottom first-principle theoretical description of a cold-atom apparatus; as such, this optimization is largely empirical but guided by theory and modeling.

Traditional optimization methods operate by sampling from and navigating the cost landscape in an attempt to locate the (global) minimum. Many contemporary approaches now frame this as a machine learning (ML) task, with the objective transformed into "learning" a high-quality model of the underlying physical reality, from which extrema can be accurately identified. Gaussian processes (GP), a type of Bayesian optimization, is one such approach, while this has already yielded promising results in cold-atom experiments,[36,37,39] many challenges remain. For instance, these optimization methods suffer from the curse of dimensionality and scale poorly with increasing parameter-count in two ways: (1) Even in principle, the required number of samples rapidly increases (thereby, increasing the requisite computational resources for Bayesian inference) and (2) in conjunction with experimental noise, large-parameter count GP models can often fail to converge. Furthermore, similar to most ML approaches, the performance of these methods depends strongly on the choice of hyperparameters,[51] and the mechanism (if any) of in-loop hyperparameter tuning as the cost function landscape is explored.

We developed a multi-stage optimization procedure to circumvent the curse of dimensionality. This procedure begins by manual scanning each parameter to identify an initial value and bounds for the following optimization. Next, these parameters are grouped into (potentially) intuitively related sets—e.g., "laser cooling" or "ODT evaporation"—of about ten parameters and optimized with GP. We then use a stochastic sampling technique to compute the parameter correlation matrix, allowing us to quantitatively collect parameters into sets with strong intra-set correlations and weak inter-set correlations, which are then separately optimized.

Introducing user-defined algebraic constraints to this optimization task is a further challenge. This is commonly implemented by adding a penalty term to the cost function.[37,52] While effective, this technique requires careful tuning of the relative strength of the penalty with respect to the remainder of the cost function: an excessive penalty can cause the optimizer to diverge, while too little will be ignored. Instead, we constrained the optimization task by including a Lagrange multiplier in the cost function. When constraining the total evaporation time, such a method enables the production of high-quality BECs with progressively shorter experimental durations.

In the remainder of this section, we describe separate strategies to effectively optimize large parameter sets (Sec. IV A) and incorporate algebraic constraints (Sec. IV C).

A. Parameter grouping

Our focus is on ML optimization tasks where experimental data acquisition rate is low (limiting the number of samples) and computational requirements scale poorly with the dimension of the parameter space. In general, this situation is hopeless and giving up is likely the most rational response. Experimentalists are not known for such sensible behavior, and besides, in laboratory settings even a highly non-trivial cost function landscape may contain some hidden structure. In our case, this structure is encoded in correlations between control parameters: parameters that are highly correlated in the neighborhood of one critical point (local minima, maxima, or saddle point) are likely to be correlated near another. This converts a hopeless situation into a hopeful three-step process of: identifying at least one critical point, obtaining the parameter's correlation matrix at that location, and optimizing parameters in sets with strong correlations.

During the setup stage of any real experiment, every control parameter is likely to have been studied at least once; for each parameter, this provides working information regarding its nominal optimal operating point, and overall range of utility. Taken together, this set of nominal optimal operating points identifies a critical point in parameter space with zero gradient of the fitness function along all directions, therefore, suitable to explore the correlation matrix. In practice, we employed several rounds of human-guided manual tuning and GP optimization to identify the initial operating point, denoted by the vector $\mathbf{x}^{(0)} \equiv \{\Delta x_0^{(0)}, \Delta x_1^{(0)}, \cdots\}$ in parameter space.

We then randomly sample the cost function at points $\mathbf{x}^{(n)}$ in the neighborhood of the initial operating point, i.e., small $\Delta x_i^{(n)} \equiv x_i - x_i^{(0)}$, and fit a second-order polynomial,

$$C(\boldsymbol{\Delta x}) \approx \mathcal{A} + \sum_i \mathcal{B}_i \Delta x_i + \sum_{i,j} \Delta x_i \mathcal{C}_{i,j} \Delta x_j \tag{1}$$

to the resulting data. Here, $\mathcal{A}$ measures the nominal value of the cost function; $\mathcal{B}_i$ is the local gradient [ideally zero, at $x_i^{(0)}$]; and $\mathcal{C}_{i,j}$ is the desired covariance matrix. The diagonal elements of $\mathcal{C}_{i,j}$ quantify the individual parameter sensitivities, and the off-diagonal elements describe cross correlations. Because the parameter sensitivities are completely arbitrary (given by whatever units happen to be used in the experimental control system), we quantify parameter correlations in terms of squared normalized correlation matrix $\bar{\mathcal{C}}_{i,j}^2 \equiv \mathcal{C}_{i,j}^2/(\mathcal{C}_{i,i}\mathcal{C}_{j,j})$ whose diagonal elements have been scaled to unity.

Figure 5 illustrates $\bar{\mathcal{C}}_{i,j}^2$ for two sets of parameters taken from two very different parts of the experimental sequence: the first set includes parameters from the laser cooling stages, while the other contains parameters from the hybrid trap transfer stage (stage IV

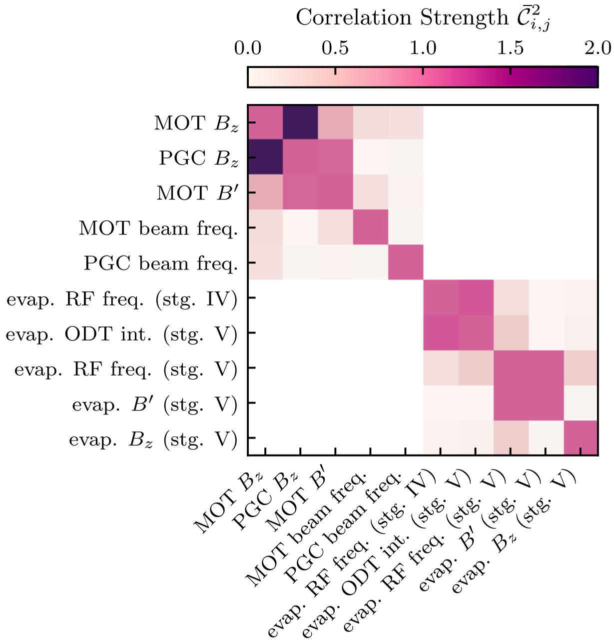

FIG. 5. Control parameter correlations. These parameters include those used for laser cooling and transfer into the hybrid trap (stage IV and V in Fig. 4). The strength of correlations is represented by the color of the off-diagonal terms. Each row and column is normalized by the value of the diagonal term (see the text).

and V in Fig. 4). As perhaps intuitively expected, only parameters within these sets are correlated. However, even within these sets, specific parameters exhibit differing degrees of correlation.

1. Laser cooling parameters

The bias and gradient magnetic fields during the MOT loading and sub-Doppler cooling stages are strongly correlated; this arises from the slow response time of the coils. By contrast, the laser detunings during the MOT loading and sub-Doppler cooling stages control very different processes and are weakly correlated.

2. Hybrid trap transfer parameters

The magnetic trap gradient and rf frequencies are highly correlated as their interplay controls the evaporation efficiency. Similarly, the rf frequency and the ODT intensity are strongly correlated because the rf frequency sets the temperature of the cloud which the ODT must then be deep enough to capture.

In the end, experimentally identified correlations such as these are used to partition the whole parameter set into strongly correlated subsets that are optimized together.

B. Bayesian optimization

We employ a GP-based Bayesian optimizer, which is particularly effective in settings where sampling the cost function is expensive, the number of sampled points is restricted, and experimental noise is present.

At every optimization, a GP optimizer encodes a probability distribution over the set of all potential cost functions $\{f(\mathbf{x})\}$, where the probability density of a specific cost function is given by a (functional) normal distribution. The task of the GP model is to identify the best choice of the distribution's $\mathbf{x}$-dependent mean $\mu^{(n)}(\mathbf{x})$ and width $\sigma^{(n)}(\mathbf{x})$ given the information available at iteration n. In practice, this problem is still intractable, so to reduce the search space GP models, employ correlation length hyperparameters that describe the range in every direction of parameter space over which functions are likely to exhibit correlations. At each iteration, a GP model generally uses the most likely function $C^{(n)}_{\text{GP}}(\mathbf{x}) = \mu^{(n)}(\mathbf{x})$ as the estimated cost function.

Given a GP model at iteration n, a numerical optimization algorithm determines the next parameter vector $\mathbf{x}^{(n+1)}$ to be sampled. One common choice is the upper confidence bound criterion: $\mathbf{x}^{(n+1)} = \arg\min\{\mu(\mathbf{x}^{(n)}) + \alpha\sigma(\mathbf{x}^{(n)})\}$, where the hyperparmeter α balances pursuing the minimum and exploring the unknown region of the model and is dynamically regulated during the optimization procedure. The potential optimal $\mathbf{x}^{(n+1)}$ is experimentally measured, yielding $C_m(\mathbf{x}^{(n+1)})$, the GP parameters are updated, and the process repeats with a new parameter vector. This iterative procedure continues until a termination criterion is met, such as a predefined number of iterations or convergence to an optimal parameter set.

We employ the open-source package "Machine-Learning Online Optimization Package" (`M-LOOP`)[36] to implement the GP optimization and interface it with our `labscript`-based[53] experiment control system. `M-LOOP` is seeded with an initial dataset obtained by experimentally sampling the cost at parameter vectors $\mathbf{x}^{(n)}$. The set of vectors might be selected at random, with a different optimizer such as differential evolution (our choice), or even with a previous GP model. The measured points help update GP model's hyperparameters. The GP optimizer allows us to efficiently explore the multiparameter space by prioritizing regions with higher uncertainty, unlike regular grid scans that uniformly sample the space. This allows faster convergence to optimal parameters compared to the conventional scan-based approach in fewer runs.

We note that the samples collected during GP optimization are insufficient to accurately extract the correlation matrix in the region surrounding the cost extrema; this is due to the low sample count required for convergence.

C. Constrained optimization

Our constrained optimization method is built upon the standard Lagrange multiplier (LM) technique. While this approach is analytically straightforward, its implementation in a data-sparse setting with experimental noise requires special consideration.

For the case of a single constraint $f(\mathbf{x}) = 0$, the LM method first introduces a new cost function $C_{\text{LM}}(\mathbf{x}) = C(\mathbf{x}) + \lambda f(\mathbf{x})$, where the coefficient λ is the LM. As usual, local minima of C_{LM} occur at points of zero gradient,

$$\nabla_{\mathbf{x}} C_{\text{LM}}(\mathbf{x}) = \nabla_{\mathbf{x}} C(\mathbf{x}) + \lambda \nabla_{\mathbf{x}} f(\mathbf{x}) = 0. \quad (2)$$

In conjunction with the constraint $f(\mathbf{x}) = 0$, this yields $n+1$ coupled equations: sufficient to resolve the $n = \dim(\mathbf{x})$ components of $\mathbf{x}$ along with the LM λ. The gradient of the constraint $\nabla_{\mathbf{x}} f(\mathbf{x})$ can be directly evaluated analytically since the constraint is known. In contrast, $\nabla_{\mathbf{x}} C(\mathbf{x})$ is an experimental quantity, which can be estimated

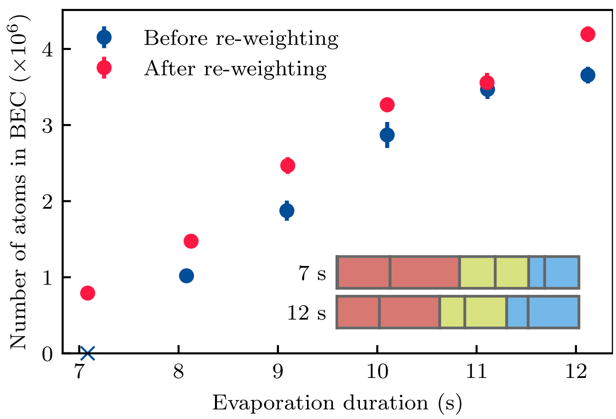

FIG. 6. Evaporation sequence performance. The blue/red dots mark the performance of the initial sequence before/after re-weighting, respectively. The blue cross at 7 s indicates that the sequence before re-weighting utterly failed in producing a BEC. The two bars in the inset indicate the percentage of time spent in each stage for the optimal 7 and 12 s sequence. The red, green, and blue regions mark the rf evaporation stage, the decompression/transfer stage, and the hybrid trap evaporation stage, respectively.

either by line scans or more efficiently via random sampling in the neighborhood of $\mathbf{x}$.

The strategy for updating the parameter vector $\mathbf{x}^{(n)}$ at iteration n to $\mathbf{x}^{(n+1)}$ proceeds as follows. First, we estimate the cost function and its gradient at $\mathbf{x}^{(n)}$ using a sampling strategy as discussed in Sec. IV A above, and obtain the gradient $\nabla_{\mathbf{x}} f(\mathbf{x}^{(n)})$ from a linear fit to the data. We then employ a constrained gradient descent to update $\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} + \beta \mathbf{P}_f \nabla_{\mathbf{x}} C(\mathbf{x}^{(n)})$ with convergence parameter β and $\mathbf{P}_f$ projects onto the surface where $f(\mathbf{x}) = 0$. Separately, we solve the n linear equations in Eq. (2), each with a separate λ_i as the unknown quantity. These λ_i's will be equal only when $\mathbf{x}^{(n)}$ is an exceptional point of the constrained optimization problem. The optimization, therefore, terminates when $|\lambda_i - \lambda_j| < \epsilon$ for all i, j.

We demonstrate this procedure by optimizing the duration of each evaporation stage while constraining their sum: the total evaporation time T. Without the constraint, the optimal evaporation sequence would balance improved thermalization (favoring longer durations) and vacuum-limited lifetime (favoring shorter durations). Constraining the total duration, therefore, strategically exchanges time among stages in a way that optimizes BEC production.

For our specific linear constraint $f(t_1, \ldots, t_n) = \sum_i t_i - T = 0$, Eq. (2) gives $\lambda_i = -\partial_{t_i} C(t_1, \ldots, t_n)$. The gradient descent process can thus be reframed as a re-weighting process in which the duration of stages with small λ_i are reduced, while the duration of stages with large λ_i are increased. This process is rapidly convergent, usually within three iterations; we tune β at each step to ensure $\sum_j |t_j^{(n)} - t_j^{(n+1)}| = 0.6\text{s}$ and terminate when $|\lambda_i - \lambda_j|$ is smaller than the largest uncertainties in all λ_i's. In practice, we perform a final stage of GP optimization (with durations) after completing the

TABLE II. List of 53 parameters that were machine-optimized as part of this study.

Experiment stage	Parameter	Description
3D MOT loading	I_{MOT}	$F = 2 \rightarrow F' = 3$ cooling intensity
	δ_{MOT}	$F = 2 \rightarrow F' = 3$ cooling detuning
	$B_{x,y,z}$	Bias fields
	B'	Quadrupole magnetic field gradient
PGC	I_{PGC}	$F = 2 \rightarrow F' = 3$ cooling intensity
	δ_{PGC}	$F = 2 \rightarrow F' = 3$ cooling detuning
	$B_{x,y,z}$	Bias fields
	t_{PGC}	PGC duration
Optical pumping	I_{depump}	$F = 2 \rightarrow F' = 2$ depump intensity
	δ_{depump}	$F = 2 \rightarrow F' = 2$ depump detuning
	I_{pump}	$F = 1 \rightarrow F' = 1$ optical pumping intensity
	$B_{x,y,z}$	Bias fields
Magnetic trap compress	B'_{capture}	Capture quadrupole magnetic field gradient
	B'_0	Final compressed quadrupole magnetic field gradient
	t_{compress}	Compression duration
	$f_{\text{rf},0}$	Initial RF frequency
	$I_{\text{ODT},0}$	Initial ODT power
Evaporation stage $j = 1, \ldots, 6$	$f_{\text{rf},j}$	Ending RF frequency
	$B_{[x,y,z],j}$	Ending bias fields
	$I_{\text{ODT},j}$	Ending ODT power
	B'_j	Ending magnetic quadrupole field gradient
	t_j	Time duration

constrained optimization process; however, at best, this yields a marginal improvement.

Figure 6 shows performance improvement of this process (red) compared to simply proportional scaling the stage durations (blue) for sequences with constrained evaporation durations from 7 to 12 s. The bar-graphs in Fig. 6 show how this process fractionally redistributes time between the different stages. As the evaporation duration is reduced, the fraction of time allocated to hybrid trap evaporation (blue) is reduced and redistributed to the rf evaporation stage (red).

We fully automated this process using our `labscript`-based experiment control system, with supporting code available online.[54]

V. CONCLUSION

Here, we described an apparatus that produces large sodium BECs in a hybrid trap and describe efficiency gains achieved via parameter grouping, GP optimization, and LM-based constrained optimization. This integrated approach strikes a practical balance between performance of the optimization system while still benefiting from correlations between parameters. This enabled our BEC apparatus to outperform existing sodium BEC systems in terms of BEC production efficiency as a function of evaporation time (see Fig. 1). We also note that improvements in laser cooling, such as D1 gray molasses,[34] could boost the phase space density prior to evaporation and circumvent the need for large-volume magnetic traps (and their attendant slow evaporation) for the production of large BECs.

Our automated time re-weighting protocol, based on constrained optimization, efficiently redistributes the evaporation sequence under fixed time constraints, significantly enhancing BEC quality by shifting resources from less effective to more impactful stages. In the results presented here, we minimized the sequence duration with fixed MOT parameters, arguing that the MOT and evaporation stages are largely decoupled. In reality, these stages are correlated (albeit weakly); this correlation can be leveraged to even more effectively reduce the total sequence duration. For example, increasing the initial number of laser cooled atoms increases the collision rate in the magnetic trap, allowing for more rapid evaporation. So, in such a situation, our LM technique may well sacrifice evaporation time for added MOT loading time.

ACKNOWLEDGMENTS

The authors thank K. T. Hoang and R. Shrestha for carefully reading the manuscript, as well as J. V. Porto and S. Subhankar for initial discussions. This work was partially supported by the National Institute of Standards and Technology and the National Science Foundation through the Quantum Leap Challenge Institute for Robust Quantum Simulation (Grant No. OMA-2120757).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Y.G. and S.M. contributed equally to this work.

Y. Geng: Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **S. Mukherjee**: Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **S. Banik**: Formal analysis (equal); Investigation (equal); Methodology (equal); Validation (equal). **M. Gutierrez Galan**: Formal analysis (equal); Investigation (equal); Methodology (equal); Validation (equal). **M. Anderson**: Formal analysis (equal); Investigation (equal); Methodology (equal); Validation (equal). **H. Sosa-Martinez**: Formal analysis (equal); Investigation (equal); Methodology (equal); Validation (equal). **S. Eckel**: Resources (equal); Supervision (equal); Writing – review & editing (equal). **I. B. Spielman**: Conceptualization (equal); Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Writing – review & editing (equal). **G. K. Campbell**: Conceptualization (equal); Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

APPENDIX: ALL MACHINE-OPTIMIZED PARAMETERS

All parameters optimized by the Bayesian optimizer are listed in Table II.

REFERENCES

[1] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell, "Observation of Bose-Einstein condensation in a dilute atomic vapor," Science **269**, 198–201 (1995).

[2] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle, "Bose-Einstein condensation in a gas of sodium atoms," Phys. Rev. Lett. **75**, 3969 (1995).

[3] I. Bloch, J. Dalibard, and W. Zwerger, "Many-body physics with ultracold gases," Rev. Mod. Phys. **80**, 885 (2008).

[4] I. Bloch, J. Dalibard, and S. Nascimbène, "Quantum simulations with ultracold quantum gases," Nat. Phys. **8**, 267–276 (2012).

[5] F. Schäfer, T. Fukuhara, S. Sugawa, Y. Takasu, and Y. Takahashi, "Tools for quantum simulation with ultracold atoms in optical lattices," Nat. Rev. Phys. **2**, 411–425 (2020).

[6] H. Ritsch, P. Domokos, F. Brennecke, and T. Esslinger, "Cold atoms in cavity-generated dynamical optical potentials," Rev. Mod. Phys. **85**, 553 (2013).

[7] M. A. Cazalilla, R. Citro, T. Giamarchi, E. Orignac, and M. Rigol, "One dimensional bosons: From condensed matter systems to ultracold gases," Rev. Mod. Phys. **83**, 1405 (2011).

[8] T. Langen, R. Geiger, and J. Schmiedmayer, "Ultracold atoms out of equilibrium," Annu. Rev. Condens. Matter Phys. **6**, 201–217 (2015).

[9] N. Navon, R. P. Smith, and Z. Hadzibabic, "Quantum gases in optical boxes," Nat. Phys. **17**, 1334–1341 (2021).

[10] L. Amico, D. Anderson, M. Boshier, J.-P. Brantut, L.-C. Kwek, A. Minguzzi, and W. von Klitzing, "Colloquium: Atomtronic circuits: From many-body physics to quantum technologies," Rev. Mod. Phys. **94**, 041001 (2022).

[11] A. Tononi and L. Salasnich, "Low-dimensional quantum gases in curved geometries," Nat. Rev. Phys. **5**, 398–406 (2023).
[12] O. Lahav, A. Itah, A. Blumkin, C. Gordon, S. Rinott, A. Zayats, and J. Steinhauer, "Realization of a sonic black hole analog in a Bose–Einstein condensate," Phys. Rev. Lett. **105**, 240401 (2010).
[13] J. Steinhauer, "Observation of self-amplifying hawking radiation in an analogue black-hole laser," Nat. Phys. **10**, 864–869 (2014).
[14] J. Steinhauer, "Observation of quantum hawking radiation and its entanglement in an analogue black hole," Nat. Phys. **12**, 959–965 (2016).
[15] S. Eckel, A. Kumar, T. Jacobson, I. B. Spielman, and G. K. Campbell, "A rapidly expanding Bose–Einstein condensate: An expanding universe in the lab," Phys. Rev. X **8**, 021021 (2018).
[16] J. R. Muñoz de Nova, K. Golubkov, V. I. Kolobov, and J. Steinhauer, "Observation of thermal hawking radiation and its temperature in an analogue black hole," Nature **569**, 688–691 (2019).
[17] V. I. Kolobov, K. Golubkov, J. R. Muñoz de Nova, and J. Steinhauer, "Observation of stationary spontaneous hawking radiation and the time evolution of an analogue black hole," Nat. Phys. **17**, 362–367 (2021).
[18] S. Banik, M. G. Galan, H. Sosa-Martinez, M. Anderson *et al.*, "Accurate determination of hubble attenuation and amplification in expanding and contracting cold-atom universes," Phys. Rev. Lett. **128**, 090401 (2022).
[19] C. Viermann, M. Sparn, N. Liebster, M. Hans, E. Kath, Á. Parra-López, M. Tolosa-Simeón, N. Sánchez-Kuntz, T. Haas, H. Strobel *et al.*, "Quantum field simulator for dynamics in curved spacetime," Nature **611**, 260–264 (2022).
[20] A. Zenesini, A. Berti, R. Cominotti, C. Rogora, I. G. Moss, T. P. Billam, I. Carusotto, G. Lamporesi, A. Recati, and G. Ferrari, "False vacuum decay via bubble formation in ferromagnetic superfluids," Nat. Phys. **20**, 558 (2024).
[21] G. Santarelli, C. Audoin, A. Makdissi, P. Laurent, G. J. Dick, and A. Clairon, "Frequency stability degradation of an oscillator slaved to a periodically interrogated atomic resonator," IEEE Trans. Ultrason. Ferroelectrics Freq. Control **45**, 887–894 (1998).
[22] K. M. R. van der Stam, E. D. van Ooijen, R. Meppelink, J. M. Vogels, and P. Van der Straten, "Large atom number Bose–Einstein condensate of sodium," Rev. Sci. Instrum. **78**, 013102 (2007).
[23] J.-y. Choi, M.-S. Heo, and Y.-i. Shin, "Experimental apparatus for producing large ^{23}Na Bose-Einstein condensates," J. Korean Phys. Soc. **59**, 211–216 (2011).
[24] E. Mimoun, L. D. Sarlo, D. Jacob, J. Dalibard, and F. Gerbier, "Fast production of ultracold sodium gases using light-induced desorption and optical trapping," Phys. Rev. A **81**, 023631 (2010).
[25] J. Jiang, L. Zhao, M. Webb, N. Jiang, H. Yang, and Y. Liu, "Simple and efficient all-optical production of spinor condensates," Phys. Rev. A **88**, 033620 (2013).
[26] W. Liu, N. Zheng, X. Wang, J. Xu, Y. Li, V. B. Sovkov, P. Li, Y. Fu, J. Wu, J. Ma, L. Xiao, and S. Jia, "Fast, simple, all-optical production of sodium spinor condensates," J. Phys. B: At., Mol. Opt. Phys. **54**, 155501 (2021).
[27] A. Farolfi, "Spin dynamics in two-component Bose–Einstein condensates," Ph.D. thesis, Università degli studi di Trento, 2021.
[28] G. Colzi, E. Fava, M. Barbiero, C. Mordini, G. Lamporesi, and G. Ferrari, "Production of large Bose-Einstein condensates in a magnetic-shield-compatible hybrid trap," Phys. Rev. A **97**, 053625 (2018).
[29] E. W. Streed, A. P. Chikkatur, T. L. Gustavson, M. Boyd, Y. Torii, D. Schneble, G. K. Campbell, D. E. Pritchard, and W. Ketterle, "Large atom number Bose–Einstein condensate machines," Rev. Sci. Instrum. **77**, 023106 (2006).
[30] Z. Shi, Z. Li, P. Wang, K. Sadiq Nawaz, L. Chen, Z. Meng, L. Huang, and J. Zhang, "Production of ^{23}Na Bose–Einstein condensates in the $F = 2$ state using D_2 gray molasses," J. Opt. Soc. Am. B **38**, 1229–1234 (2021).
[31] W. Hänsel, P. Hommelhoff, T. Hänsch, and J. Reichel, "Bose–Einstein condensation on a microelectronic chip," Nature **413**, 498–501 (2001).
[32] D. M. Farkas, E. A. Salim, and J. Ramirez-Serrano, "Production of rubidium Bose–Einstein condensates at a 1 Hz rate," arXiv:1403.4641 (2014).
[33] Y.-J. Lin, A. R. Perry, R. L. Compton, I. B. Spielman, and J. V. Porto, "Rapid production of ^{87}Rb Bose–Einstein condensates in a combined magnetic and optical potential," Phys. Rev. A **79**, 063631 (2009).
[34] G. Colzi, G. Durastante, E. Fava, S. Serafini, G. Lamporesi, and G. Ferrari, "Sub-Doppler cooling of sodium atoms in gray molasses," Phys. Rev. A **93**, 023421 (2016).
[35] M.-S. Heo, J.-y. Choi, and Y.-i. Shin, "Fast production of large ^{23}Na Bose–Einstein condensates in an optically plugged magnetic quadrupole trap," Phys. Rev. A **83**, 013622 (2011).
[36] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn *et al.*, "Fast machine-learning online optimization of ultra-cold-atom experiments," Sci. Rep. **6**, 25890 (2016).
[37] Z. Vendeiro, J. Ramette, A. Rudelis, M. Chong, J. Sinclair, L. Stewart, A. Urvoy, and V. Vuletić, "Machine-learning-accelerated Bose–Einstein condensation," Phys. Rev. Res. **4**, 043216 (2022).
[38] E. T. Davletov, V. V. Tsyganok, V. A. Khlebnikov, D. A. Pershin, D. V. Shaykin, and A. V. Akimov, "Machine learning for achieving Bose-Einstein condensation of thulium atoms," Phys. Rev. A **102**, 011302 (2020).
[39] R. Roussel, A. L. Edelen, T. Boltz, D. Kennedy, Z. Zhang, F. Ji, X. Huang, D. Ratner, A. S. Garcia, C. Xu, J. Kaiser, A. F. Pousa, A. Eichler, J. O. Lübsen, N. M. Isenberg, Y. Gao, N. Kuklev, J. Martinez, B. Mustapha, V. Kain, C. Mayes, W. Lin, S. M. Liuzzo, J. St John, M. J. . Streeter, R. Lehe, and W. Neiswanger, "Bayesian optimization algorithms for accelerator physics," Phys. Rev. Accel. Beams **27**, 084801 (2024).
[40] S. Banik, "Analogue cosmology experiments with sodium Bose–Einstein condensates," Ph.D. thesis, University of Maryland, 2021.
[41] M. Gutierrez Galan, "Bose Einstein condensates for analogue cosmology experiments," Ph.D. thesis, University of Maryland, 2021.
[42] G. Lamporesi, S. Donadello, S. Serafini, and G. Ferrari, "Compact high-flux source of cold sodium atoms," Rev. Sci. Instrum. **84**, 063102 (2013).
[43] Certain equipment, instruments, software, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement of any product or service by NIST, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.
[44] T. G. Tiecke, S. D. Gensemer, A. Ludewig, and J. T. M. Walraven, "High-flux two-dimensional magneto-optical-trap source for cold lithium atoms," Phys. Rev. A **80**, 013409 (2009).
[45] A. Gozzini, F. Mango, J. H. Xu, G. Alzetta, F. Maccarrone, and R. A. Bernheim, "Light-induced ejection of alkali atoms in polysiloxane coated cells," Il Nuovo Cimento D **15**, 709–722 (1993).
[46] G. Telles, T. Ishikawa, M. Gibbs, and C. Raman, "Light-induced atomic desorption for loading a sodium magneto-optical trap," Phys. Rev. A **81**, 032710 (2010).
[47] W. Ketterle, K. B. Davis, M. A. Joffe, A. Martin, and D. E. Pritchard, "High densities of cold atoms in a dark spontaneous-force optical trap," Phys. Rev. Lett. **70**, 2253–2256 (1993).
[48] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf, "Observation of atoms laser cooled below the Doppler limit," Phys. Rev. Lett. **61**, 169 (1988).
[49] D. M. Stamper-Kurn, H.-J. Miesner, A. P. Chikkatur, S. Inouye, J. Stenger, and W. Ketterle, "Reversible formation of a Bose–Einstein condensate," Phys. Rev. Lett. **81**, 2194–2197 (1998).
[50] C. J. Pethick and H. Smith, *Bose–Einstein Condensation in Dilute Gases* (Cambridge University Press, 2001).
[51] P. T. Sivaprasad, F. Mai, T. Vogels, M. Jaggi, and F. Fleuret, "Optimizer benchmarking needs to account for hyperparameter tuning," in *Proceedings of the 37th International Conference on Machine Learning, Proceedings of the Machine Learning Research*, edited by H. D. III and A. Singh (PMLR, 2020), Vol. 119, pp. 9036–9045.
[52] A. J. Barker, H. Style, K. Luksch, S. Sunami, D. Garrick, F. Hill, C. J. Foot, and E. Bentine, "Applying machine learning optimization methods to the production of a quantum gas," Mach. Learn.: Sci. Technol. **1**, 015007 (2020).
[53] P. T. Starkey, C. J. Billington, S. P. Johnstone, M. Jasperse, K. Helmerson, L. D. Turner, and R. P. Anderson, "A scripted control system for autonomous hardware-timed experiments," Rev. Sci. Instrum. **84**, 085111 (2013).
[54] See https://github.com/JQIamo/labscript-time-reweight-optimization for "JQIamo / labscript-time-reweight-optimization" (last accessed 29 August 2025).