Demonstration of a programmable optical lattice atom interferometer

Catie LeDesma,[1] Kendall Mehling,[1] Jieqiu Shao,[2] John Drew Wilson,[1] Penina Axelrad,[3] Marco Nicotra,[2] Dana Z. Anderson,[1,4] and Murray Holland[1,*]

[1]*JILA & Department of Physics, University of Colorado Boulder, Boulder, Colorado 80309, USA*
[2]*College of Engineering and Applied Science, University of Colorado Boulder, Boulder, Colorado 80309, USA*
[3]*Department of Aerospace Engineering, University of Colorado Boulder, Boulder, Colorado 80303, USA*
[4]*Infleqtion, 3030 Sterling Circle, Boulder, Colorado 80301, USA*

(Received 24 July 2023; accepted 2 October 2024; published 8 November 2024)

Performing interferometry in an optical lattice formed by standing waves of light offers potential advantages over its free-space equivalents since the atoms can be confined and manipulated by the optical potential. We demonstrate such an interferometer in a one-dimensional lattice and show the ability to control the atoms by imaging and reconstructing the wave function at many stages during its cycle. An acceleration signal is applied, and the resulting performance is seen to be close to the optimum possible for the time-space area enclosed according to quantum theory. Our methodology of machine design enables the sensor to be reconfigurable on the fly, and when scaled up, offers the potential to make state-of-the art inertial and gravitational sensors that will have a wide range of potential applications.

DOI: 10.1103/PhysRevResearch.6.043120

I. INTRODUCTION

In this paper, we utilize machine learning and optimal control methods to manipulate atoms confined in an optical lattice to sense accelerations. Atom interferometry is well established as a means of measuring inertial forces with exquisite sensitivity [1,2]. Scientifically compelling endeavors that can be advanced through precision interferometry often present constraints or environments that are challenging for typical atom interferometric systems to accommodate. These include low orbit monitoring of the gravity field of Earth [3–5], and spaceborne searches for dark matter [6,7], both of which are constrained by size and weight limitations and present harsh vibrational and thermal environments. In another precision measurement context, timekeeping systems that confine atoms in optical lattices have achieved stunning levels of precision—displaying a fractional frequency uncertainty of 3.5×10^{-19} [8] for shallow lattices, and 7.6×10^{-21} [9] for deep lattices. Timekeeping experiments have established that an optical lattice can provide a pristine environment for precision metrology [10], even while the forces imposed on atoms correspond to the range of tens to hundreds of g's, where g is the acceleration due to the gravity of Earth. From a practical standpoint, therefore, optical lattices can be used to confine and manipulate atoms in the face of a dynamically harsh environment [11]; the question then arises whether this system can be used for matter-wave interferometry. The answer lies in machine-learning methods, which can be used to discover how an atomic wave function associated with a lattice can be manipulated by changing the relative phase of the optical lattice to achieve interferometric measurement precision [12,13], see Fig. 1(a).

Optical lattices have been used to enhance the performance of atom interferometric inertial sensors in a variety of ways, such as to demonstrate the high sensitivity of a sensor based on Bloch oscillations [14], to capitalize on long atomic coherence times by holding separated atoms in place [15], and as a means to impose large momentum transfer to atoms [16,17]. In this paper, we demonstrate an interferometry sequence where the atoms are confined within the optical lattice potential during the entirety of the measurement sequence. Unlike other optical lattice based sensors, our device efficiently transfers atoms into the conduction band of the optical periodic potential, enabling the atoms to enclose an extended enclosed area for enhanced sensitivity to inertial signals. We call this method Bloch-band interferometry (BBI) due to the fact that, during the free propagation of our sensing sequence, the lattice remains stationary, and the atomic wave function can be described as being in a superposition of Bloch states. BBI combines the approach of the shaken-lattice interferometer [18,19], where the lattice is dynamically shaken to enable atoms to move through the valence band of the lattice via dynamic tunneling [20], and that of lattice-based interferometry, where atoms are held stationary in a lattice during phase accumulation. In BBI, the lattice is only shaken during application of interferometery protocols, i.e., beamsplitters and mirrors, and is held stationary during propagation steps, allowing atoms to traverse the lattice as essentially free particles rather than through tunneling. The power of this optical lattice approach combined with optimal control methods to generate a nearly arbitrary quantum state was illustrated in Ref. [21]. Motivated by these earlier works, here, we

*Contact author: murray.holland@colorado.edu

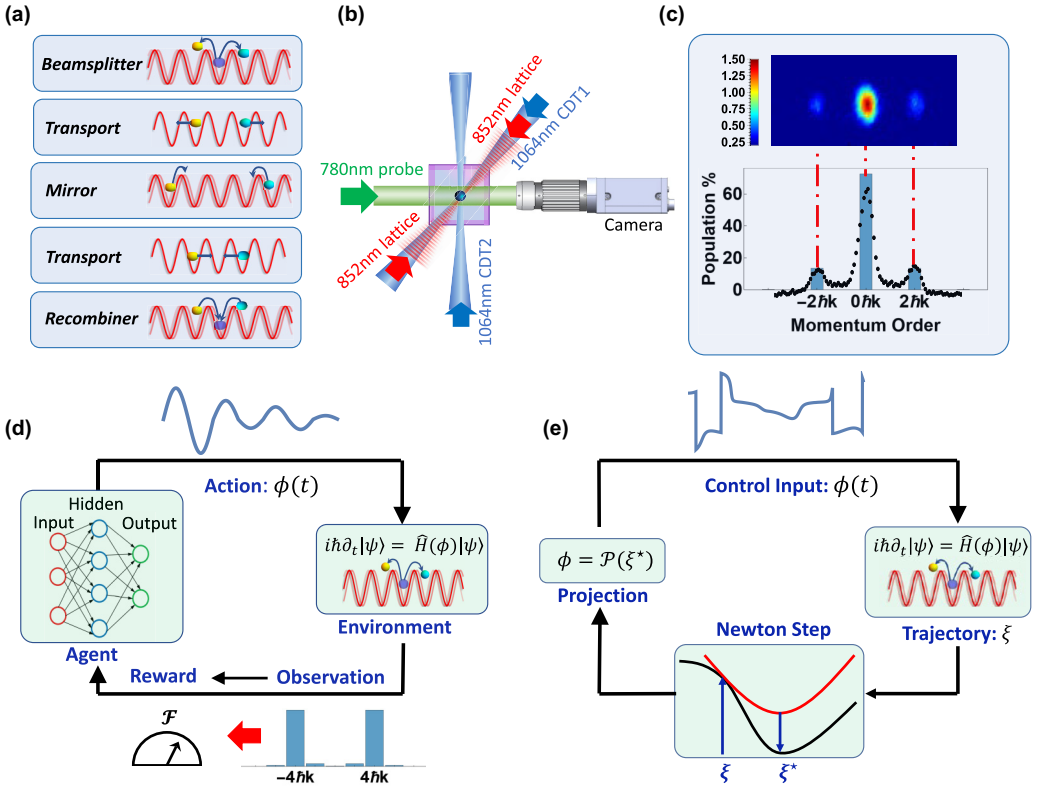

FIG. 1. Experimental system: (a) The lattice is shaken to implement the interferometer components. (b) Condensates are created in two crossed 1064 nm dipole laser beams (blue) that intersect inside a science cell (purple). The atoms are then loaded into a one-dimensional 852 nm optical lattice (red), whose position can be altered by applying frequency shifts to two acousto-optic modulators (AOMs; not shown). After time of flight (TOF), an absorption image of the cloud is taken by a 780 nm probe laser (green). (c) A TOF absorption image of the optical density (OD) of the cloud after lattice load. Also shown is the integrated momentum distribution (lower) overlaid with the theoretical histogram anticipated for the lowest Bloch eigenstate at that lattice depth. Machine-design methods: (d) Reinforcement learning, where an agent (neural network) chooses an action, the environment responds to the action, and a reward is given as feedback based on the observed state. (e) Quantum optimal control using the PRojection Operator Newton method for Trajectory Optimization (PRONTO) algorithm [22]. For each iteration, given the control input $\phi(t)$, PRONTO first computes the trajectory ξ based on the Schrödinger equation, then the algorithm uses the second-order approximation of the cost, which is a function of ξ, to find an updated trajectory $\xi^\star$

demonstrate a machine-designed optical lattice atom interferometer and evaluate its performance as a sensor.

II. EXPERIMENTAL PROCEDURE

As shown in Fig. 1(b), we produce 2×10^4 Bose-condensed ^{87}Rb atoms via all-optical evaporation of atoms in a 1064 nm crossed dipole trap (CDT) [23]. At the end of evaporation, the dipole beams are maintained at their final setpoint to offset gravity while the optical lattice beams are adiabatically ramped on, leading to the transfer of condensed atoms into the ground-state wave function with the momentum distribution shown in Fig. 1(c). The lattice is formed by two counterpropagating 852 nm laser beams, whose intensity and phase are each controlled by independent acousto-optic modulators (AOMs) before entering the science cell. The atoms are subject to an optical potential at point $\mathbf{r}$ and time t:

$$V(\mathbf{r}, t) = V_L(\mathbf{r}_\perp)\cos[2\mathbf{k}\cdot\mathbf{r} + \phi(t)] + V_D(\mathbf{r}), \qquad (1)$$

where wave vector $\mathbf{k}$ points along the lattice direction (with magnitude $k = 2\pi/\lambda$ given by the wavelength λ of the light), and $V_L(\mathbf{r}_\perp)$ is the depth of the lattice with $\mathbf{k}\cdot\mathbf{r}_\perp = 0$. We refer to $\phi(t)$ as a control or shaking function. The potential from the CDT, $V_D(\mathbf{r})$, which remains present during the shaking, gives rise to an ellipsoidal trap [24] with frequencies of 36.5, 132.3, and 144.8 Hz, defined along the principal axes and measured via parametric heating [25]. Total atom numbers are calibrated by means of time of flight expansion and fitting to the Thomas-Fermi scaling ansatz [26,27]. To measure the momentum distribution, all confining beams are rapidly extinguished, and the atoms are allowed to expand via TOF for 15 ms. Normalized atom numbers for each momentum state are extracted from the integrated optical density of the probe image. To accurately calibrate our lattice depth, we either perform Kapitza-Dirac diffraction [28,29] and fit with theory or measure the frequency of momentum state oscillations following a small sudden shift of the optical lattice position [30].

A. Reinforcement learning and optimal control methods

The problem of quantum design has been recast in recent years by the developments in machine-learning methods,

which can be extremely effective at finding efficient strategies for accomplishing complex tasks [31]. These computer-based approaches are capable of being decoupled from human intuition and consequently have often proven effective at uncovering solution spaces that have never been previously explored. When naïve brute force searches are prohibitive, sophisticated algorithms, such as those used in reinforcement learning, have been able to perform at a level that exceeds human capacity, as demonstrated by the strength of computers playing games such as chess and Go [32–34]. The quantum domain is strongly connected with this paradigm due to the analogous exponential complexity underpinning quantum evolution when the system is composed of many constituents [35–41]. Examples are finding protocols for quantum communication [42], improving quantum sensors [43], and engineering quantum currents [44].

As schematically illustrated in Figs. 1(d) and 1(e), we employ two distinct machine-design approaches: reinforcement learning (RL) and quantum optimal control (QOC), to find lattice position profiles, i.e., control functions, to realize desired momentum transformations. In this paper, all control functions are produced by simulating the one-dimensional Schrödinger equation for an infinite lattice, with constraints on the lattice depth and total integration time. RL involves the application of an agent which is implemented using a deep neural network. This agent repeatedly proposes a sequence of actions chosen from a predetermined set to discover a sequence that leads to a high terminal reward based on the quantum state fidelity. QOC adopts a different strategy that assumes a Hamiltonian model of the wave function evolution and computes the control sequence that minimizes a given cost function.

The two approaches differ in that RL iterates through many shaking sequences and typically arrives at a different solution in every run, while for the same control seed and design parameters, QOC will always return the same optimal control shaking function. It is not particularly useful to compare the two methods per se, as they operate in different design spaces. In the case of QOC, the model is fully specified, and the algorithm is deterministic, while in the case of RL, the learning is model free. In model-free learning, it is possible, at least in principle, to perform closed-loop learning on an experimental system whose dynamical equations are completely unknown. Notwithstanding these considerations, both learning methods have proven to be outstanding at finding high-quality solutions in our design landscape. Acceptable simulated solutions for beamsplitters and mirrors typically achieve >97% fidelity and are often >99%.

B. Beamsplitter design

We apply both RL and QOC methods to implement offline machine-designed solutions for interferometer components. The first component in the interferometer sequence is an atomic beamsplitter, which places each atom in a superposition of traveling in both directions at the same time. The protocol that we implement was previously developed in Ref. [12] using a $10E_R$ deep lattice, which is the lattice depth used exclusively for all the experimental sequences presented here. The lattice depth is expressed in terms of the recoil energy $E_R = \hbar^2 k^2/2m$, where m is the atomic mass of ^{87}Rb. With this depth, we realize an approximate beamsplitter by targeting the third excited Bloch state $|n = 3, q = 0\rangle \equiv |3\rangle$, where n is the band index and q is the quasimomentum of the periodic lattice. This target state was chosen since it is simultaneously an eigenstate of the lattice and is primarily composed of the two $\pm 4\hbar k$ momentum states with equal weight (∼47% in each). An additional motivation for selecting this target state is that the atoms are excited into the conduction band, enabling them to cover large distances during the transport segment of the interferometer sequence.

Figures 2(a)–2(d) show the outcomes following the application of the machine-designed beamsplitter shaking function. Although the shaking functions produced by RL and QOC are dramatically different, the measured fidelities to the target beamsplitter are both >95% [21]. Both RL and QOC can produce beamsplitters having total shaking times much shorter than than the recoil period, $\tau_R = 2\pi/\omega_R \sim 316\,\mu$s. We note that these are very short component time scales, of the order of the vibrational period of the atoms in the potential wells of the lattice.

C. Mirror design

Designing a shaking function that serves as a mirror is a substantially different problem than the beamsplitter case. For the latter, the design seeks to transform an initial wave function into a final target wave function. In the mirror case, the target is a unitary operation [12] rather than a state. Mathematically, this means a mirror unitary U_M is designed to perform the following transformation:

$$U_M(\alpha\,|p_0\rangle + \beta\,|-p_0\rangle) \Longrightarrow \beta\,|p_0\rangle + \alpha\,|-p_0\rangle\,, \qquad (2)$$

where $|\pm p_0\rangle = \frac{1}{\sqrt{2}}(|3\rangle \pm |4\rangle)$ are the approximately pure $\pm 4\hbar k$ momentum eigenstates, and channel fidelity [45] is used as the metric of performance. Using channel fidelity restricts the unitary optimization to the subspace spanned by $\{|p_0\rangle\,, |-p_0\rangle\}$, which is the only part of the quantum space that is important for this mirror.

To validate mirror performance, we first design a shaking sequence that produces a 100/0 beamsplitter (sending all the atoms to the $|p_0\rangle$ state) or a 0/100 beamsplitter (sending all the atoms to the $|-p_0\rangle$ state). This allows us to observe the quality of the momentum reflection directly. Figure 3 presents the experimental evolution from the ground state to a 100/0 beamsplitter state and the subsequent mirror action.

D. The Michelson interferometer

Equipped with beamsplitter and mirror protocols, a full interferometer sequence is stitched together, effectively splitting, propagating, mirroring, propagating, and finally, recombining for phase readout, as shown in Figs. 4(a)–4(c). We impose time-reversal symmetry so that only the first half of the interferometer sequence need be designed.

Following the entire interferometer sequence, the atoms are largely returned to the ground state of the lattice. This is anticipated if the components are implemented accurately since, ideally, in the absence of an externally applied acceleration signal, the initial state should be perfectly recovered at the end of the sequence. The fact that we observe this, as shown

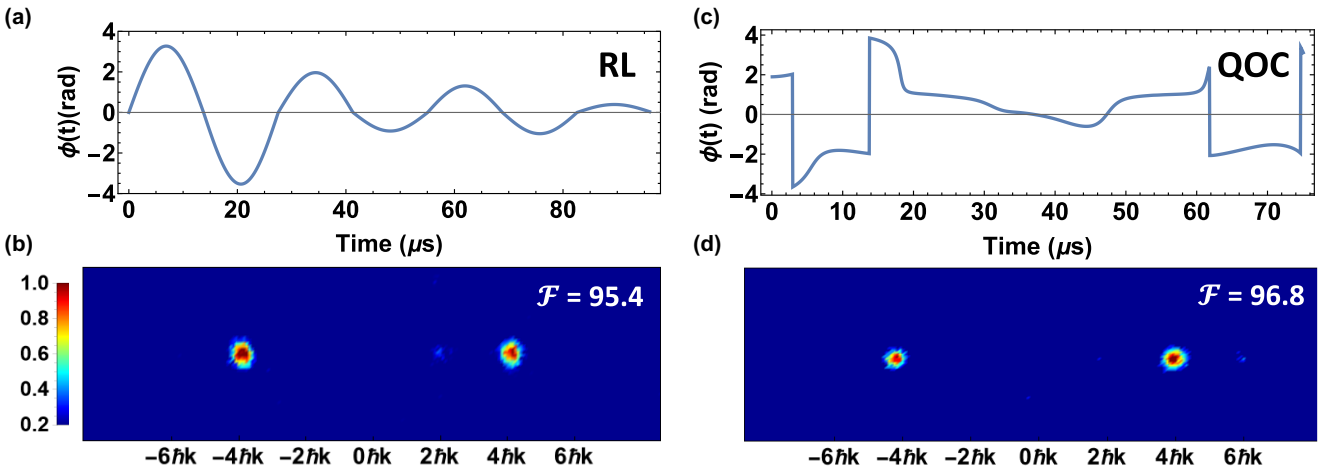

FIG. 2. Beamsplitter sequences: (a) and (b) Sequence learned by reinforcement learning (RL) with the resulting time of flight images from experiment. Similar measured momentum distributions for quantum optimal control (QOC) are depicted in (c) and (d). The experimental fidelities for these sequences to the target solutions $\mathcal{F}$, as defined in Ref. [21], are as indicated. In (b) and (d), color bars denote the optical density.

in Fig. 4, indicates not only that the devices are operating very close to what was intended in terms of both phase and amplitude wave function evolution but also that any effects not modeled by the design equations do not adversely affect solutions on these time scales.

E. Interferometer response to acceleration

In a standard interferometer, the output is dependent upon an applied acceleration. The sensitivity to this acceleration is proportional to the enclosed spacetime area. The shaken lattice approach allows us to adjust the enclosed area by

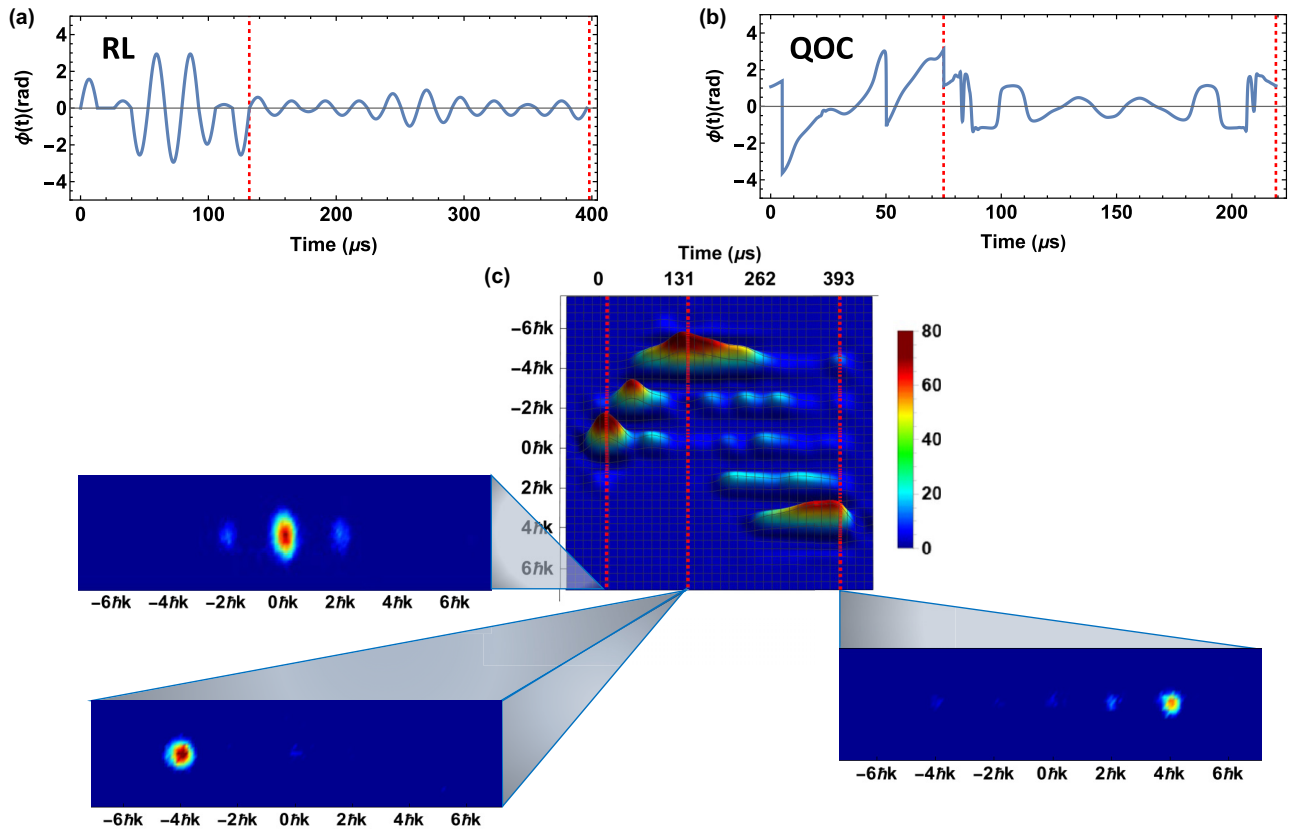

FIG. 3. Mirror sequences: Shaking sequences for the 100/0 beamsplitter and mirror for (a) quantum optical control (QOC) and (b) reinforcement learning (RL). Red dashed lines indicate the start and end of the mirror sequence. (c) The cascaded time of flight images for RL, as seen in experiment, with popouts showing the initial Bose-Einstein condensate (BEC) loaded into the lattice, then after the 100/0 beamsplitter, and finally, after the mirror is applied. The momentum transfer from $-4\hbar k$ to $4\hbar k$ by the mirror is apparent. Analogous results are observed for sending the atoms along the other path, and/or using QOC control functions.

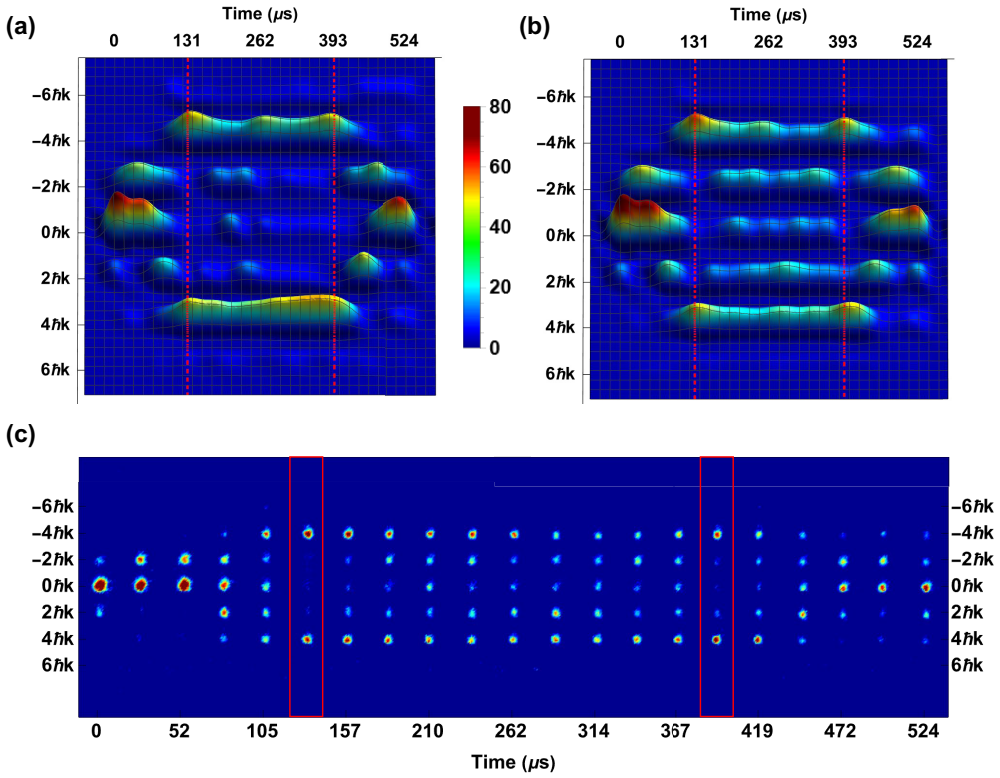

FIG. 4. Interferometer sequences: (a) Experimental momentum probability distributions as observed from time of flight expansion and cascaded together to give a representation of the interferometer evolution. The color bar is expressed in units of normalized population percentage, and red lines indicate the boundaries between the beamsplitter, mirror, and recombiner sequences, which are sandwiched together with no propagation delays. The design problem specifies the desired wave function only on these boundary lines. (b) Same depiction as (a) but for the theoretical quantum evolution. (c) Raw experimental TOF images stitched together. Note that the observed asymmetry between positive and negative momenta is anticipated since shaking solutions immediately break the inversion symmetry by initially shifting the lattice in one direction.

incorporating transport stages between the interferometer components, see Fig. 1(a). During the transport stages, the lattice position is held fixed in space. Here, we allow the atoms to propagate for 100 μs in each direction. We keep this time relatively short to ensure that any potential effects of transverse diffusion of the atoms, walk-off from the CDT intersection, and atom-atom interactions will not significantly impact the sensor performance. An acceleration is imposed onto the atoms via a frequency chirp applied to one of the lattice AOMs. Although both design methods produce comparable results, we display only the QOC results for the sake of brevity, with the control function shown in Fig. 5(a). The resulting momentum state interference fringes as a function of acceleration are shown in Fig. 5(b). The experimental momentum populations measured over the applied acceleration range compare well with the theory presented in Fig. 5(c). Note that the only fitting parameter is a small overall shift of the zero value in the applied accelerations which is accounted for in the data presented in Figs. 5(d) and 5(e). This shift is caused by a slight tilt ($<1^\circ$) in our lattice beams, which induces an acceleration offset.

There are several significant aspects of the resulting fringes, as illustrated in Fig. 5(d), that we draw attention to. The output of our accelerometer is the momentum state population fractions $p \in \{-6\hbar k, -4\hbar k, \ldots, 6\hbar k\}$ viewed after recombination and TOF, rather than the interference signal of two recombined clouds as in conventional Bragg interferometry. The fraction of the total population in these seven momentum states $P(p|a)$ provides a unique fingerprint for each acceleration a, which allows the value of the acceleration to be determined from the measured momentum populations using statistical methods. The populations are not symmetric about the applied accelerations, enabling the discrimination of the direction of the applied signal. The fringe periodicity, which determines the dynamic range of the device, can easily be tuned by simply changing the transport time, with longer times giving higher sensitivity and smaller bandwidth and shorter times giving lower sensitivity and higher bandwidth. This ability to modify the system performance metrics at will is particularly attractive for applications in which the measurement scenario is dynamic.

F. Sensitivity

The performance of the atom accelerometer can be characterized by the the Jensen-Shannon (JS) divergence illustrated in Fig. 5(e). The JS divergence quantifies the degree of consistency between the simulation and the experimental measurements [46]. An unbiased estimator for the measured acceleration is the maximum likelihood, which corresponds to the minimum of the JS divergence [47] at each measured data point and is shown as white circles. The degree to which

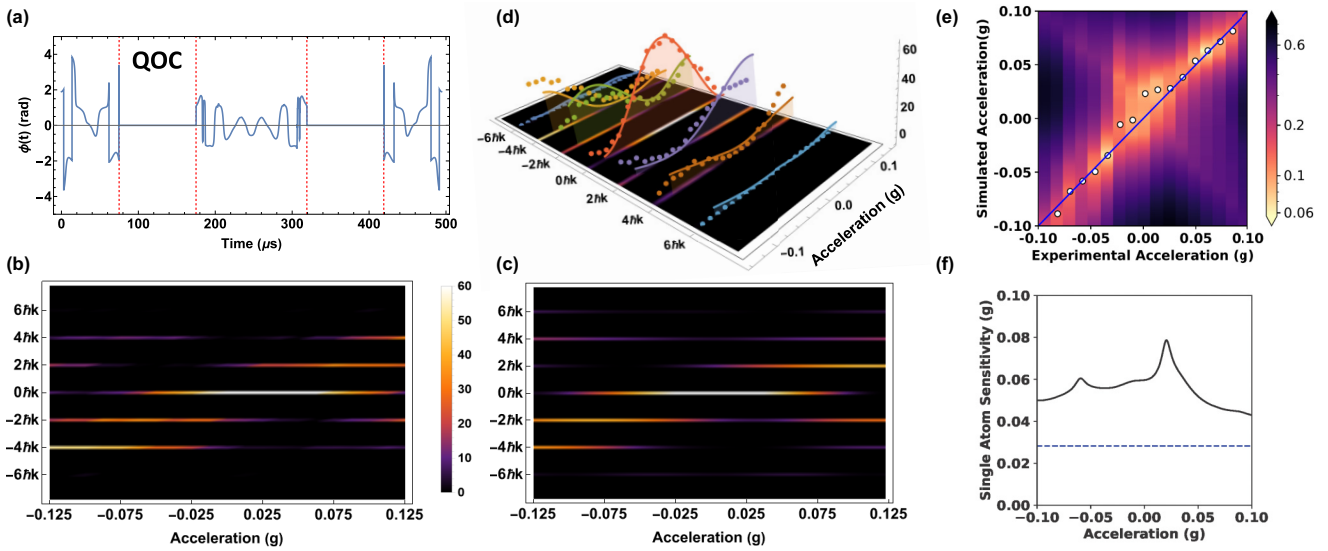

FIG. 5. Accelerometer measurements: (a) The quantum optical control (QOC) interferometer control sequence, where red lines separate interferometer components. Full interferometer sequences were conducted over a set of equally spaced acceleration values from $-0.125g$ to $+0.125g$, and time of flight momentum populations were measured and stitched together to generate the contoured scan depicted in (b). The expected result from the theory calculation is shown in (c). Fringes for each momentum component are depicted in (d). The experimental data points (dots) compare well with the anticipated theory (solid lines). The theory lines shown are not fits to the data and contain no free parameters with the exception of the overall shift in the zero acceleration. Analysis of bias offset and sensitivity: (e) Jensen-Shannon divergence shown on a logarithmic scale. Low values ~ 0 correspond to almost perfect coincidence of the measured distribution with theory (light regions), and high values ~ 1 correspond to poor coincidence (dark regions). The minimum at each recorded data distribution is the maximum likelihood estimator for acceleration (white circles). (f) Single-atom sensitivity (smaller values are more sensitive) from theory for this shaking sequence compared with the ideal (instantaneous component and perfect momentum splitting) limit with the same total time (blue dashed line).

the calibrated points lie in proximity to the diagonal line is a measure of sensitivity. To quantify this, we compute the single-atom sensitivity $1/\sqrt{I(a)}$, where $I(a)$ is the classical Fisher information:

$$I(a) = \sum_{p\in\{-6\hbar k,\ldots,6\hbar k\}} \frac{1}{P(p|a)}\left[\frac{\partial P(p|a)}{\partial a}\right]^2, \tag{3}$$

which is shown in Fig. 5(f). For reference, we bound this by the ideal single-atom sensitivity that would result if the beamsplitter and mirror components took zero time and the device operated in free space with exactly $\pm 4\hbar k$ splitting, thereby giving a perfect two-arm interferometer.

III. DISCUSSION AND CONCLUSIONS

In this paper, we have demonstrated the application of advanced machine-design methods to find a class of quantum design solutions for phase modulation, that is, the shaking of an optical lattice. We have shown the ability to produce a variety of beamsplitters (50/50, 100/0, and 0/100), mirrors, and recombiners, all of which have been realized experimentally. When the designed components were cascaded, the resulting interferometer protocols were analyzed through the application of a lattice acceleration and the subsequent measurement of the output diffraction pattern.

The interferometer demonstrated here has a round-trip time of $<500\,\mu$s. The transport time in the lattice ($100\,\mu$s) implies that the atoms only move on the order of $2.5\,\mu$m. This is important because the sensitivity of an accelerometer scales with the area enclosed in spacetime. If the transport time were increased from $100\,\mu$s to, say, 0.1 s, the atomic wave packets would be displaced by ~ 2.5 mm; this would increase sensitivity by a factor of order $(10^3)^2$. According to Fig. 5(f), this implies a possible single-atom sensitivity of $\approx 6\times 10^{-8}g$. This value is not a sensitivity limit since each fringe can be more precisely determined by statistical averaging using multiple atoms, multiple passes, and multiple trials [48]. The anticipated performance at the shot noise limit can, in general, be computed from the single-atom sensitivity by dividing by a factor of $\sqrt{N}$ for N independent measurements. Also note that what we present here is a classical limit, so that entangling and squeezing the atomic quantum state could potentially improve the sensitivity further.

It should also be noted that, because we use machine learning for design, protocols can be constructed to mitigate the adverse effects of laser noise and atom interactions [49]. This can be done by including each of these effects in the reward function used during learning, a process that involves application of the multiparameter quantum Fisher information matrix [50]. This leads to design protocols that are extremely sensitive to acceleration, while being simultaneously insensitive to imperfections, such as lattice depth and/or atom number fluctuations, commonly classified as nuisance parameters. The ability to design around imperfections is a feature that can be used to make optical-lattice interferometry attractive for deployment in environments where robustness is needed.

In summary, in this paper, we have realized one of a class of atom interferometers that are real-time reconfigurable, have a compact form factor, and the potential to extend the

application domain of atom interferometry. For more details, see the Supplemental Material [51].

ACKNOWLEDGMENTS

We would like to thank Liang-Ying Chih, Malcolm Boshier, Ceren Uzun, Katarzyna Krzyżanowska, and Mantas Naris for many helpful discussions. This paper was supported by NSF OMA Grant No. 1936303, NSF PHY Grant No. PHY 2317149, NSF OMA Grant No. 2016244, and NSF PHY Grant No. 2207963, and by a Sponsored Research Agreement between the University of Colorado Boulder and Infleqtion.

[1] P. Asenbaum, C. Overstreet, M. Kim, J. Curti, and M. A. Kasevich, Atom-interferometric test of the equivalence principle at the 10^{-12} level, Phys. Rev. Lett. **125**, 191101 (2020).

[2] J. Lee, R. Ding, J. Christensen, R. R. Rosenthal, A. Ison, D. P. Gillund, D. Bossert, K. H. Fuerschbach, W. Kindel, P. S. Finnegan *et al.*, A compact cold-atom interferometer with a high data-rate grating magneto-optical trap and a photonic-integrated-circuit-compatible laser system, Nat. Commun. **13**, 5131 (2022).

[3] G. M. Tino, A. Bassi, G. Bianco, K. Bongs, P. Bouyer, L. Cacciapuoti, S. Capozziello, X. Chen, M. L. Chiofalo, A. Derevianko *et al.*, Sage: A proposal for a space atomic gravity explorer, Eur. Phys. J. D **73**, 228 (2019).

[4] M. D. Lachmann, H. Ahlers, D. Becker, A. N. Dinkelaker, J. Grosse, O. Hellmig, H. Müntinga, V. Schkolnik, S. T. Seidel, T. Wendrich *et al.*, Ultracold atom interferometry in space, Nat. Commun. **12**, 1317 (2021).

[5] A. Trimeche, B. Battelier, D. Becker, A. Bertoldi, P. Bouyer, C. Braxmaier, E. Charron, R. Corgier, M. Cornelius, K. Douch *et al.*, Concept study and preliminary design of a cold atom interferometer for space gravity gradiometry, Class. Quantum Gravity **36**, 215004 (2019).

[6] P. Hamilton, M. Jaffe, P. Haslinger, Q. Simmons, H. Müller, and J. Khoury, Atom-interferometry constraints on dark energy, Science **349**, 849 (2015).

[7] Y.-D. Tsai, J. Eby, and M. S. Safronova, Direct detection of ultralight dark matter bound to the sun with space quantum sensors, Nat. Astron. **7**, 113 (2023).

[8] K. Kim, A. Aeppli, T. Bothwell, and J. Ye, Evaluation of lattice light shift at low 10^{-19} uncertainty for a shallow lattice Sr optical clock, Phys. Rev. Lett. **130**, 113203 (2023).

[9] T. Bothwell, C. J. Kennedy, A. Aeppli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, Resolving the gravitational redshift across a millimetre-scale atomic sample, Nature (London) **602**, 420 (2022).

[10] H. Katori, Optical lattice clocks and quantum metrology, Nat. Photonics **5**, 203 (2011).

[11] K. Bongs, M. Holynski, J. Vovrosh, P. Bouyer, G. Condon, E. Rasel, C. Schubert, W. P. Schleich, and A. Roura, Taking atom interferometric quantum sensors from the laboratory to real-world applications, Nat. Rev. Phys. **1**, 731 (2019).

[12] L.-Y. Chih and M. Holland, Reinforcement-learning-based matter-wave interferometer in a shaken optical lattice, Phys. Rev. Res. **3**, 033279 (2021).

[13] L.-Y. Chih, D. Z. Anderson, and M. Holland, Reinforcement learning for rotation sensing with ultracold atoms in an optical lattice, arXiv:2212.14473.

[14] Z. Pagel, W. Zhong, R. H. Parker, C. T. Olund, N. Y. Yao, and H. Müller, Symmetric Bloch oscillations of matter waves, Phys. Rev. A **102**, 053312 (2020).

[15] C. D. Panda, M. Tao, J. Egelhoff, M. Ceja, V. Xu, and H. Müller, Quantum metrology by one-minute interrogation of a coherent atomic spatial superposition, Nat. Phys. **20**, 1234 (2024).

[16] K. D. Nelson, C. D. Fertig, P. Hamilton, J. M. Brown, B. Estey, H. Müller, and R. L. Compton, Guided matter wave inertial sensing in a miniature physics package, Appl. Phys. Lett. **116**, 234002 (2020).

[17] M. Gebbe, J.-N. Siemß, M. Gersemann, H. Müntinga, S. Herrmann, C. Lämmerzahl, H. Ahlers, N. Gaaloul, C. Schubert, K. Hammerer *et al.*, Twin-lattice atom interferometry, Nat. Commun. **12**, 2544 (2021).

[18] S. Pötting, M. Cramer, and P. Meystre, Momentum-state engineering and control in Bose-Einstein condensates, Phys. Rev. A **64**, 063613 (2001).

[19] C. A. Weidner, H. Yu, R. Kosloff, and D. Z. Anderson, Atom interferometry using a shaken optical lattice, Phys. Rev. A **95**, 043624 (2017).

[20] H. Lignier, C. Sias, D. Ciampini, Y. Singh, A. Zenesini, O. Morsch, and E. Arimondo, Dynamical control of matter-wave tunneling in periodic potentials, Phys. Rev. Lett. **99**, 220403 (2007).

[21] N. Dupont, G. Chatelain, L. Gabardos, M. Arnal, J. Billy, B. Peaudecerf, D. Sugny, and D. Guéry-Odelin, Quantum state control of a Bose-Einstein condensate in an optical lattice, PRX Quantum **2**, 040303 (2021).

[22] J. Shao, J. Combes, J. Hauser, and M. M. Nicotra, Projection-operator-based Newton method for the trajectory optimization of closed quantum systems, Phys. Rev. A **105**, 032605 (2022).

[23] M. D. Barrett, J. A. Sauer, and M. S. Chapman, All-optical formation of an atomic Bose-Einstein condensate, Phys. Rev. Lett. **87**, 010404 (2001).

[24] G.-B. Liao, K.-S. Wu, C.-Y. Shih, Y.-H. Cheng, L.-A. Sun, Y.-J. Lin, and M.-S. Chang, Optimization of a crossed optical dipole trap for loading and confining laser-cooled atoms, J. Opt. Soc. Am. B **34**, 869 (2017).

[25] D. Wineland, P. Ekstrom, and H. Dehmelt, Monoelectron oscillator, Phys. Rev. Lett. **31**, 1279 (1973).

[26] Y. Castin and R. Dum, Bose-Einstein condensates in time dependent traps, Phys. Rev. Lett. **77**, 5315 (1996).

[27] F. Dalfovo, C. Minniti, S. Stringari, and L. Pitaevskii, Nonlinear dynamics of a Bose condensed gas, Phys. Lett. A **227**, 259 (1997).

[28] P. L. Gould, G. A. Ruff, and D. E. Pritchard, Diffraction of atoms by light: The near-resonant Kapitza-Dirac effect, Phys. Rev. Lett. **56**, 827 (1986).

[29] S. B. Cahn, A. Kumarakrishnan, U. Shim, T. Sleator, P. R. Berman, and B. Dubetsky, Time-domain de Broglie wave interferometry, Phys. Rev. Lett. **79**, 784 (1997).

[30] C. Cabrera-Gutiérrez, E. Michon, V. Brunaud, T. Kawalec, A. Fortun, M. Arnal, J. Billy, and D. Guéry-Odelin, Robust

calibration of an optical-lattice depth based on a phase shift, Phys. Rev. A **97**, 043617 (2018).

[31] 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).

[32] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, Mastering the game of Go with deep neural networks and tree search, Nature (London) **529**, 484 (2016).

[33] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton *et al.*, Mastering the game of Go without human knowledge, Nature (London) **550**, 354 (2017).

[34] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel *et al.*, A general reinforcement learning algorithm that masters chess, Shogi, and Go through self-play, Science **362**, 1140 (2018).

[35] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Reinforcement learning in different phases of quantum control, Phys. Rev. X **8**, 031086 (2018).

[36] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, Reinforcement learning with neural networks for quantum feedback, Phys. Rev. X **8**, 031084 (2018).

[37] H. P. Nautrup, N. Delfosse, V. Dunjko, H. J. Briegel, and N. Friis, Optimizing quantum error correction codes with reinforcement learning, Quantum **3**, 215 (2019).

[38] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, Universal quantum control through deep reinforcement learning, npj Quantum Inf. **5**, 33 (2019).

[39] M. Bukov, Reinforcement learning for autonomous preparation of Floquet-engineered states: Inverting the quantum Kapitza oscillator, Phys. Rev. B **98**, 224305 (2018).

[40] J. Mackeprang, D. B. R. Dasari, and J. Wrachtrup, A reinforcement learning approach for quantum state engineering, Quantum Mach. Intell. **2**, 5 (2020).

[41] A. A. Melnikov, P. Sekatski, and N. Sangouard, Setting up experimental Bell tests with reinforcement learning, Phys. Rev. Lett. **125**, 160401 (2020).

[42] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel, Machine learning for long-distance quantum communication, PRX Quantum **1**, 010301 (2020).

[43] J. Schuff, L. J. Fiderer, and D. Braun, Improving the dynamics of quantum sensors with reinforcement learning, New J. Phys. **22**, 035001 (2020).

[44] T. Haug, R. Dumke, L.-C. Kwek, C. Miniatura, and L. Amico, Machine-learning engineering of quantum currents, Phys. Rev. Res. **3**, 013034 (2021).

[45] L. H. Pedersen, N. M. Møller, and K. Mølmer, Fidelity of quantum operations, Phys. Lett. A **367**, 47 (2007).

[46] J. Lin and S. Wong, A new directed divergence measure and its characterization, Int. J. Gene. Syst. **17**, 73 (1990).

[47] I. Csiszár and F. Matus, Information projections revisited, IEEE Trans. Inf. Theory **49**, 1474 (2003).

[48] H. Kim, K. Krzyzanowska, K. C. Henderson, C. Ryu, E. Timmermans, and M. Boshier, One second interrogation time in a 200 round-trip waveguide atom interferometer, arXiv:2201.11888.

[49] L.-Y. Chih, Machine-learning-based design of quantum systems for extreme sensing, Ph.D. thesis, CU Boulder, 2022.

[50] J. T. Reilly, J. D. Wilson, S. B. Jäger, C. Wilson, and M. J. Holland, Optimal generators for quantum sensing, Phys. Rev. Lett. **131**, 150802 (2023).

[51] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevResearch.6.043120 for additional details concerning the experimental apparatus, the methods used to develop machine-learned protocols for interferometry, and further material on distance metrics used for validating the sensitivity of the device.