

Reinforcement learning for rotation sensing with ultracold atoms in an optical lattice

Liang-Ying Chih¹ and Murray Holland^{1*}

JILA and Department of Physics, University of Colorado, 440 UCB, Boulder, Colorado 80309, USA



(Received 29 December 2022; accepted 1 November 2024; published 25 November 2024)

In this paper, we investigate a design approach of reinforcement learning to engineer a gyroscope in an optical lattice for the inertial sensing of rotations. Our methodology is not based on traditional atom interferometry, that is, splitting, reflecting, and recombining wavefunction components. Instead, the learning agent is assigned the task of generating lattice shaking sequences that optimize the sensitivity of the gyroscope to rotational signals in an end-to-end design philosophy. What results is an interference device that is completely distinct from the familiar Mach-Zehnder-type interferometer. For the same total interrogation time, the end-to-end design leads to a twentyfold improvement in sensitivity over traditional Bragg interferometry.

DOI: 10.1103/PhysRevResearch.6.043191

I. INTRODUCTION

Over the past decade, optical Michelson interferometers have achieved astounding levels of sensitivity leading to entirely new fields of gravitational wave cosmology [1,2]. Their matter-wave counterparts, in particular atom interferometers [3,4], have proven to be a principal technology for sensing [5–16], especially for the inertial measurements of accelerations, rotations, gravimetry [17], and gravity gradiometry [18]. The sensitivity that can be achieved in atom interferometry is fundamentally limited by the de Broglie wavelength of the atoms, the area that can be enclosed in space and/or space-time, and by the number of particles [19]. Improvements have primarily focused on individual components, such as larger momentum splitting and longer hold time.

Recently, however, a completely distinct approach to light pulse atom interferometry has been developed [20,21]. It involves confining atoms to an optical lattice during the entire interferometry protocol, and builds on early experiments that created beamsplitters and other components by shaking the lattice [22]. In part, the purpose of the optical lattice is to provide robustness of the system in the face of a harsh dynamical environment typical of real-world applications [23,24], including platform vibrations and temperature fluctuations and drifts. Since there is no unique protocol for generating the components of beam-splitters and mirrors in an optical lattice, it makes sense to delegate this optimization task to machine learning algorithms [21]. One “teaches” the system to carry out an interferometric task by learning how to correctly modulate the phase of the lattice. Experiments using this method can now demonstrate full interferometer sequences that cascade the machine-designed components to-

gether [25–27]. In these experiments it has been shown that the learning can be performed in simulation first and only later applied to the experiment, as demonstrated by the high-fidelity agreement of the measured data with the anticipated sensor performance.

While these works have established the key principles, the learning that was implemented was constrained by a premeditated viewpoint as to how interferometry is to take place. Traditionally, nearly every kind of interferometry, optical or matter wave, results from a sequence of wavefront splitting, reflection, and recombination, with wave propagation in between, known as the Mach-Zehnder configuration [28]. However, if the learning agent is freed from being restricted by conventional wisdom, it could potentially discover solutions that surpass what humans have explored to date [29]. This different learning goal may provide a pathway to the discovery of a revolutionary kind of atomic gyroscope that is purely optimized for specified design constraints, such as a combination of sensitivity, dynamic range, and tolerance to noise and experimental drifts.

A type of machine learning that has the potential to realize this goal is reinforcement learning [30]. Reinforcement learning is especially effective when the agent is trained in an end-to-end manner, meaning an uninterrupted shaking sequence from start to finish that is not decomposable into components such as beam splitters and mirrors. This is because reinforcement learning allows for the maximization of long-term rewards of a control protocol in situations where the optimal solution is not transparent at each individual step. Reinforcement learning has emerged as a powerful tool in quantum research and has already been applied to a variety of quantum problems, such as the generation of highly-entangled states and the development of quantum error correction protocols, thereby bringing new insights to the field [31–43].

In this paper, we present results from the application of reinforcement learning to quantum metrology, specifically to train a two-dimensional optical lattice to sense rotation with high precision and thereby embody the design objectives of a gyroscope.

*Contact author: murray.holland@colorado.edu

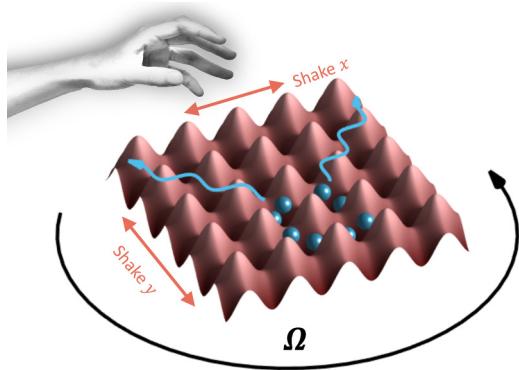


FIG. 1. A two-dimensional optical lattice potential in a rotating frame, where atoms are trapped and their dynamics can be controlled by an external agent (the hand) shaking the lattice. The blue arrows indicate possible trajectories associated with splitting of the matterwaves.

II. PHYSICAL MODEL

The system we consider is one where the atoms are confined in a two-dimensional optical lattice, as shown in Fig. 1. The lattice in each dimension can be “shaken” by varying the phases of the corresponding pairs of interfering laser beams. The entire system is described mathematically in a rotating noninertial frame by the Hamiltonian

$$\hat{H} = \frac{\hat{\mathbf{p}}^2}{2m} - \boldsymbol{\Omega} \cdot \hat{\mathbf{L}} - \sum_{\alpha \in \{x,y\}} \frac{V_\alpha}{2} \cos(2k\hat{a} + \phi_\alpha), \quad (1)$$

where m is the mass of a single atom with coordinate $\hat{\mathbf{r}} = \{\hat{x}, \hat{y}\}$ and momentum $\hat{\mathbf{p}}$, k is the laser wavenumber, ϕ_α ’s are the phase differences between the two counterpropagating lasers in the x or y directions, and V_α ’s are the corresponding lattice potential strengths. The term $\boldsymbol{\Omega} \cdot \hat{\mathbf{L}}$ describes the rotational kinetic energy of the system, with $\hat{\mathbf{L}} = \hat{\mathbf{r}} \times \hat{\mathbf{p}}$ the

angular momentum, and $\boldsymbol{\Omega}$ the angular velocity. It is the unknown magnitude, Ω , that is the metrological parameter that we wish to measure with high accuracy. For simplicity, we limit the discussion here to the case of dilute atom clouds, meaning atom-atom interactions do not adversely affect the detection of the rotation signal.

Before investigating end-to-end solutions, we first derive a conventional two-path Sagnac matter-wave interferometer, motivated by the fiberoptic gyroscope [44], but implemented with atoms [45–47]. This can be constructed in the lattice using the previously developed beam-splitting and reflecting protocols that were based on components optimized by reinforcement learning [21]. The splitting protocol was demonstrated to transfer the ground state to an approximate superposition of the $|\pm 4\hbar k\rangle$ states, and the reflecting protocol to map any linear combination of the $|\pm 4\hbar k\rangle$ states to the corresponding combination of the $|\mp 4\hbar k\rangle$ states. In order to build a gyroscope from these components, as illustrated in Fig. 2(a), the sequence of operations is the following:

(i) The wavefunction is prepared in the ground state of the 2D lattice.

(ii) In the x direction, the 1D beam-splitting protocol is applied, allowing free propagation for a duration, denoted as T , and then the reflecting protocol is applied. Following this, free propagation occurs for a duration $2T$, and then the reflecting protocol is applied again, followed by free propagation for a duration T , and finally the recombining protocol.

(iii) In the y direction, the lattice operates as a conveyor belt and is smoothly accelerated to the velocity $4\hbar k/m$ according to the adiabaticity criterion, as described in the Supplemental Material (SM) [48], and then translated at constant velocity. During the halfway point of the sequence, when the two paths in the x -directions cross, the velocity of the y -lattice is decelerated through zero to $-4\hbar k/m$, so that the lattice can be translated backward at the constant velocity. The last step is to adiabatically accelerate the y lattice back to zero velocity for the final recombination.

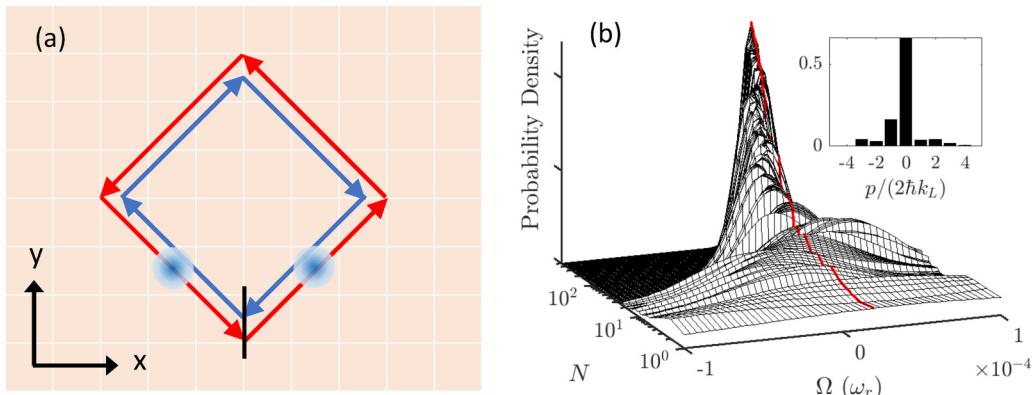


FIG. 2. (a) The clockwise trajectory (blue) and the counter-clockwise trajectory (red) taken by the wave functions in the two-path shaken lattice gyroscope. The vertical back line is where the wave functions were initially split and are eventually recombined. (b) Bayesian reconstruction of the rotation rate from the interference patterns generated by the two-path gyroscope. The main plot shows the probability distribution of the estimated rotation rate for numbers of measurements ranging from 1 to 500 and lattice depths $V_x = V_y = 10\hbar\omega_r$, where $\omega_r = \hbar k^2/(2m)$. The red curve indicates the true value of the rotation rate, which was $2 \times 10^{-4} \omega_r$. At 500 measurements, the peak of the distribution coincides with the true value, meaning that the estimation is unbiased. The inset shows the momentum distribution that results at the output for a rotation rate of $2 \times 10^{-4} \omega_r$.

Such a device operates on the principle of the Sagnac effect, which describes the phase difference, $\Delta\phi$, that accumulates between two waves that propagate in opposite directions in a loop when in a noninertial rotating frame. In general, the phase difference is proportional to both the angular velocity and to the area that the two paths enclose, A . In the case of matterwaves, $\Delta\phi = 4m\Omega \cdot A/\hbar$ [49,50]. If higher sensitivity is desired, one could effectively increase A by completing many cycles before applying the last recombination step. In practice, one must balance multiple cycles against any reduced fringe visibility arising from imperfections.

III. END-TO-END DESIGN

With respect to this reference system, we now consider the possibility for a completely end-to-end design that is not based on the decomposition into components. We implement the machine learning for this overarching design goal by employing the double-deep-Q-network approach [51], an algorithm that is particularly effective in situations where deterministic outcomes are expected and sample efficiency is critical (details described in Ref. [21]). This reinforcement learning algorithm consists of an agent that selects actions based on the observed state of the environment, and an environment that carries out the the action and generates a consequential reward based on the resulting quality of the state. The agent does not need to be exposed to the full quantum state, but only to relevant features of that state, such as the population in the discrete momentum basis in x and the average position and momentum in y .

Motivated in part by potential experimental implementation, we choose the actions to be selected from a finite set of discrete options for the time derivative of the phase, $d\phi/dt$. This represents a frequency difference between interfering laser beam pairs. The wave packet evolution that this generates is calculated by numerically solving the time-dependent Schrödinger equation, using a separation ansatz for the x and y dimensions (see the SM). In order to compute the reward, we use the classical Fisher information [52], which quantitatively measures the sensitivity. To do this, we replicate the environment as three copies of the quantum system with slightly different rotation rates, $\Omega = [\Omega_0 - d\Omega, \Omega_0, \Omega_0 + d\Omega]$, for small $d\Omega$ and with each evolving according to Eq. (1). This construction allows us to compute the derivative with respect to Ω by numerical symmetric finite differencing. The classical Fisher information is given by

$$I(\Omega) = \sum_p \frac{1}{\Pr(p|\Omega)} \left[\frac{\partial \Pr(p|\Omega)}{\partial \Omega} \right]^2, \quad (2)$$

where $\Pr(p|\Omega)$ is the probability for measuring a momentum p at the gyroscope output for the given Ω . Reinforcement learning aims to generate steps that maximize the long-term reward, that is, to find a sequence of $d\phi/dt$ steps that optimizes the Fisher information evaluated at the terminal time. The potential sensitivity is constrained by a theoretical bound (Cramér-Rao bound) for the standard deviation of the measurement,

$$\sigma_\Omega \geq 1/\sqrt{NI(\Omega)}, \quad (3)$$

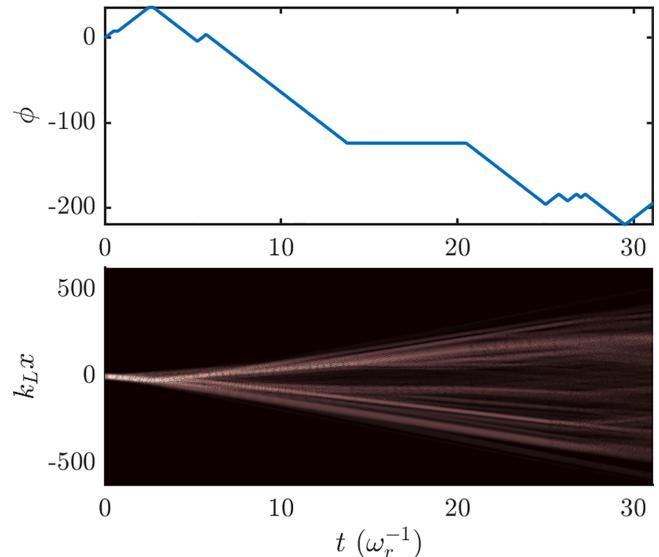


FIG. 3. Evolution of the phase function and the spatial wavefunction. The shaking protocol was generated from the reinforcement learning agent with $\partial\phi/\partial t$ as the actions in an end-to-end design. There are 125 discrete steps in this protocol. The wavefunction was initialized in the ground state of the lattice convolved with a momentum width $\sigma_p = 0.1 \hbar k$, and then evolved according to the shaking protocol given.

where N is the number of independent measurements (atoms). In this work, we have optimized for a specific choice of the angular frequency, Ω_0 , but straightforward extensions are possible, including designing for net performance over a finite domain.

IV. RESULTS

Within the context of our learning framework, we are able to obtain lattice shaking protocols that outperform the conventional two-path gyroscope. The solution varies for every instantiation due to the randomized initialization of the agent. The illustration in Fig. 3 shows one of the realized machine-learned protocols that gives high sensitivity. The pattern is reminiscent of speckle patterns that emerge from a multimode fiber (fiber specklegram), which are known to be sensitive interference detectors of inertial phase. However, in the case of a multimode fiber system, the patterns are not robust and are scrambled by temperature changes or strain on the fiber. The situation is quite different here, since the adverse noise and imperfections that may enter are different in origin and primarily common mode. Despite the nontrivial and irregular pattern, this device will measure rotation signals with high accuracy, something that we now demonstrate by simulating an example experimental measurement.

The momentum distributions that are produced by the multipath interferometer under conditions of different values of the rotation rate are shown in Fig. 4. The population of each momentum component is indicated by false color. Momentum distributions are typically directly observable in experiment by expanding the cloud under time of flight, and integrating regions of the absorption image to find the relative proportions for each momentum component. The more detailed structure

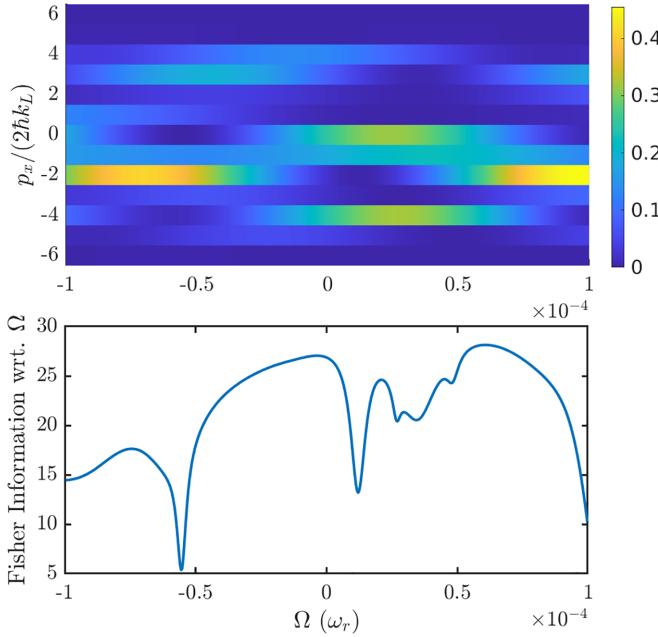


FIG. 4. Population distribution in the momentum basis and the resulting classical Fisher information for the machine-learned shaken-lattice gyroscope as a function of the rotation rate Ω . The value of the Fisher information is normalized by the Fisher information for the two-path interferometer, which has splitting of order $\Delta p = 8\hbar k$, and thus sensitivity about four times better than a conventional Bragg interferometer.

there is in the observed momentum distribution as the rotation rate is varied, the more sensitive is the interferometer. Note that each vertical slice is essentially unique, and therefore the momentum distribution is a fingerprint that allows one to infer the rotation rate without aliasing. The Fisher information calculated from the distribution is also shown in Fig. 4. The value on the y axis is the ratio between the Fisher information from this multimode interferometer and the one shown earlier based on the two-path arrangement. We optimize the Fisher information around $\Omega_0 = 0$, and one can see that the Fisher information is maximum at $\Omega = 0$, as anticipated, with large dynamic range being sacrificed for high sensitivity.

The result of our design was a machine-learned gyroscope that achieved a Fisher information of around 25 times higher than the previously investigated two-path gyroscope, meaning that the sensitivity is improved by the square root of this, or a factor of five. In other words, the interferometer is effectively as sensitive as a conventional one that is five times larger. This gain can be understood from several aspects, including population transfer into the higher momentum states (the maximum observed splitting is of order $18\hbar k$), and a larger spatial footprint, as also seen in Ref. [29]. If we make a comparison to the gyroscope with the same total interrogation time as constructed from conventional Bragg interferometry, the improvement is $4 \times 5 = 20$ -fold. The extra factor of four introduced here is primarily due to large angle splitting of the two-path gyroscope with respect to the Bragg one, i.e., $8\hbar k$ instead of $2\hbar k$, as was shown in Ref. [21]. The potential gain of 20 is remarkable and reveals the scope

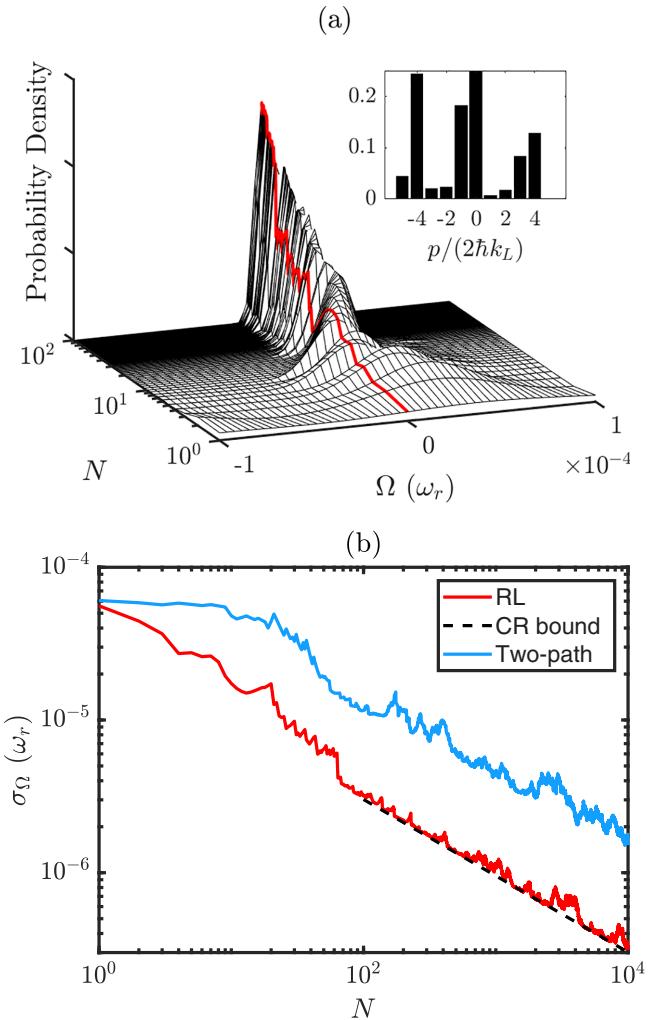


FIG. 5. (a) An example Bayesian reconstruction of the rotation rate, showing the probability distribution for sampling 1 atom to 100 atoms. The true rotation value is indicated by the red line, and the inset shows the shaken lattice interferometer output distribution for this value. (b) The standard deviation of the rotation rate as a function of the atom number (RL, red curve). Note that the behavior of the curve is subject to large sampling noise at small N , and only begins to agree with its asymptotic limit at around $N = 100$. The result at large atom number, when the central limit theorem applies and the sample average converges to the true value, is consistent with the Cramér-Rao (CR, dotted line) bound. The performance greatly surpasses the two-path interferometer shown earlier (two-path, blue curve).

for significant improvements upon current state-of-the-art devices.

To extract the rotational signal from the distribution shown in Fig. 4, we apply Bayesian reconstruction [53]. This means that we iteratively update the prior distribution for Ω from each atom measurement. The reconstruction of the rotational signal is shown in Fig. 5(a), and note that the peak coincides with the true rotation rate as indicated by the red line, thereby verifying that the estimation for Ω is unbiased. In Fig. 5(b), we plot the standard deviation of Ω , and see that the sensitivity scales inversely with the square root of the number

of atoms for up to 10^4 atoms, as expected from independent measurements. We emphasize that for large atom numbers, interactions between atoms may become increasingly important and are not considered here.

V. DISCUSSION

In conclusion, we have shown how to build a gyroscope in an optical lattice that uses demonstrated experimental methodology [25], but extended here to end-to-end design. From a broader perspective, these results show the exceptional potential of the general learning approach for quantum design, and opens up the application of the same principal machinery to a variety of quantum sensing problems with more complex landscapes, such as multiparameter estimation and

entanglement-enhanced metrology [54,55]. Furthermore, the implications of our results may go beyond this framework, and the same ideas of reinforcement learning may be applied to a variety of quantum and classical systems, such as those where the design of complex quantum circuits for algorithmic tasks is needed [56].

ACKNOWLEDGMENTS

The authors acknowledge helpful discussions with Dana Anderson, Marco Nicotra, Penina Axelrad, Catie LeDesma, and Kendall Mehling. This work was supported by the National Science Foundation (Grants No. OSI-1936303, No. PHY-2207963, No. OSI-2016244, and No. JILA-PFC-2207963).

[1] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Observation of gravitational waves from a binary black hole merger, *Phys. Rev. Lett.* **116**, 061102 (2016).

[2] F. Acernese, P. Amico, M. Al-Shourbagy *et al.*, The virgo interferometric gravitational antenna, *Optics and Lasers in Engineering* **45**, 478 (2007).

[3] C. Bordé, Atomic interferometry with internal state labelling, *Phys. Lett. A* **140**, 10 (1989).

[4] M. Kasevich and S. Chu, Atomic interferometry using stimulated raman transitions, *Phys. Rev. Lett.* **67**, 181 (1991).

[5] S. M. Dickerson, J. M. Hogan, A. Sugarbaker, D. M. S. Johnson, and M. A. Kasevich, Multiaxis inertial sensing with long-time point source atom interferometry, *Phys. Rev. Lett.* **111**, 083001 (2013).

[6] G. Raithel, A. Duspayev, B. Dash, S. C. Carrasco, M. H. Goerz, V. Vuletić, and V. S. Malinovsky, Principles of tractor atom interferometry, *Quantum Sci. Technol.* **8**, 014001 (2023).

[7] C. Schubert, S. Abend, M. Gersemann, M. Gebbe, D. Schlippe, P. Berg, and E. M. Rasel, Multi-loop atomic sagnac interferometry, *Sci. Rep.* **11**, 16121 (2021).

[8] C. Ryu and M. Boshier, Integrated coherent matter wave circuits, *New J. Phys.* **17**, 092002 (2015).

[9] T. Wilkason, M. Nantel, J. Rudolph, Y. Jiang, B. E. Garber, H. Swan, S. P. Carman, M. Abe, and J. M. Hogan, Atom interferometry with floquet atom optics, *Phys. Rev. Lett.* **129**, 183202 (2022).

[10] J. P. Cooling and J. A. Dunningham, Quantum-enhanced atomic gyroscope with tunable precision, *J. Phys. B: At., Mol. Opt. Phys.* **54**, 195502 (2021).

[11] C. D. Panda, M. Tao, J. Egelhoff, M. Ceja, V. Xu, and H. Müller, Coherence limits in lattice atom interferometry at the one-minute scale, *Nat. Phys.* **20**, 1234 (2024).

[12] C. D. Panda, M. Tao, J. Egelhoff, M. Ceja, A. Reynoso, V. Xu, and H. Müller, Probing gravity for one minute with an optical-lattice atom interferometer, *CPT and Lorentz Symmetry* (World Scientific Publishing, Singapore, 2023), pp. 136–139.

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

[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] Y.-J. Chen, A. Hansen, G. W. Hoth, E. Ivanov, B. Pelle, J. Kitching, and E. A. Donley, Single-source multiaxis cold-atom interferometer in a centimeter-scale cell, *Phys. Rev. Appl.* **12**, 014019 (2019).

[16] D. Yankelev, C. Avinadav, N. Davidson, and O. Firstenberg, Atom interferometry with thousand-fold increase in dynamic range, *Sci. Adv.* **6**, eabd0650 (2020).

[17] A. Peters, K. Y. Chung, and S. Chu, Measurement of gravitational acceleration by dropping atoms, *Nature (London)* **400**, 849 (1999).

[18] M. J. Snadden, J. M. McGuirk, P. Bouyer, K. G. Haritos, and M. A. Kasevich, Measurement of the earth's gravity gradient with an atom interferometer-based gravity gradiometer, *Phys. Rev. Lett.* **81**, 971 (1998).

[19] P. Story and C. Cohen-Tannoudji, The Feynman path integral approach to atomic interferometry. a tutorial, *J. Phys. II France* **4**, 1999 (1994).

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

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

[22] J. H. Denschlag, J. E. Simsarian, H. Häffner, C. McKenzie, A. Browaeys, D. Cho, K. Helmerson, S. L. Rolston, and W. D. Phillips, A Bose-Einstein condensate in an optical lattice, *J. Phys. B* **35**, 3095 (2002).

[23] V. Xu, M. Jaffe, C. D. Panda, S. L. Kristensen, L. W. Clark, and H. Müller, Probing gravity by holding atoms for 20 seconds, *Science* **366**, 745 (2019).

[24] 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).

[25] C. LeDesma, K. Mehling, J. Shao, J. D. Wilson, P. Axelrad, M. M. Nicotra, D. Z. Anderson, and M. Holland, Demonstration of a programmable optical lattice atom interferometer, *Phys. Rev. Res.* **6**, 043120 (2024).

[26] C. LeDesma, K. Mehling, and M. Holland, Vector atom accelerometry in an optical lattice, [arXiv:2407.04874](https://arxiv.org/abs/2407.04874) [quant-ph].

[27] C. LeDesma, K. Mehling, J. D. Wilson, M. Nicotra, and M. Holland, Universal gate set for optical lattice based atom interferometry, [arXiv:2410.17472](https://arxiv.org/abs/2410.17472) [quant-ph].

[28] L. Mach, Modifikation und Anwendung des Jamin Interferenz-Refraktometers, *Anz. Akad. Wiss. Wien math. Naturwiss.-Klasse* **28**, 223 (1891); L. Zehnder, Ein neuer interferenzrefraktor, *Z. Instrumentenkd.* **11**, 275 (1891).

[29] M. Kritsotakis, S. S. Szigeti, J. A. Dunningham, and S. A. Haine, Optimal matter-wave gravimetry, *Phys. Rev. A* **98**, 023629 (2018).

[30] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction* (MIT Press, Cambridge, MA, USA, 2018).

[31] A. A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel, Active learning machine learns to create new quantum experiments, *Proc. Natl. Acad. Sci.* **115**, 1221 (2018).

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

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

[34] M. Bukov, Reinforcement learning for autonomous preparation of Floquet-engineered states: Inverting the quantum Kapitza oscillator, *Phys. Rev. B* **98**, 224305 (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] J. Wallnöfer, A. A. Melnikov, W. Dür, and H. J. Briegel, Machine learning for long-distance quantum communication, *PRX Quantum* **1**, 010301 (2020).

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

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

[39] 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).

[40] Z. T. Wang, Y. Ashida, and M. Ueda, Deep reinforcement learning control of quantum cartpoles, *Phys. Rev. Lett.* **125**, 100401 (2020).

[41] H. Xu, J. Li, L. Liu, Y. Wang, H. Yuan, and X. Wang, Generalizable control for quantum parameter estimation through reinforcement learning, *npj Quantum Inf.* **5**, 82 (2019).

[42] X.-M. Zhang, Z. Wei, R. Asad, X.-C. Yang, and X. Wang, When does reinforcement learning stand out in quantum control? a comparative study on state preparation, *npj Quantum Inf.* **5**, 85 (2019).

[43] S. A. Haine and J. J. Hope, Machine-designed sensor to make optimal use of entanglement-generating dynamics for quantum sensing, *Phys. Rev. Lett.* **124**, 060402 (2020).

[44] V. Vali and R. W. Shorthill, Fiber ring interferometer, *Appl. Opt.* **15**, 1099 (1976).

[45] T. L. Gustavson, P. Bouyer, and M. A. Kasevich, Precision rotation measurements with an atom interferometer gyroscope, *Phys. Rev. Lett.* **78**, 2046 (1997).

[46] A. Lenef, T. D. Hammond, E. T. Smith, M. S. Chapman, R. A. Rubenstein, and D. E. Pritchard, Rotation sensing with an atom interferometer, *Phys. Rev. Lett.* **78**, 760 (1997).

[47] K. Krzyzansowska, J. Ferreras, C. Ryu, E. C. Samson, and M. Boshier, Matter wave analog of a fiber-optic gyroscope, *Phys. Rev. A* **108**, 043305 (2023).

[48] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.6.043191> for further information on the adiabaticity criterion, the reduction to an effective 1D model, and parameters used in the reinforcement learning implementation.

[49] F. Riehle, T. Kisters, A. Witte, J. Helmcke, and C. J. Bordé, Optical Ramsey spectroscopy in a rotating frame: Sagnac effect in a matter-wave interferometer, *Phys. Rev. Lett.* **67**, 177 (1991).

[50] B. Barrett, R. Geiger, I. Dutta, M. Meunier, B. Canuel, A. Gauguet, P. Bouyer, and A. Landragin, The sagnac effect: 20 years of development in matter-wave interferometry, *C. R. Phys.* **15**, 875 (2014).

[51] H. Van Hasselt, A. Guez, and D. Silver, Deep reinforcement learning with double q-learning, in *Thirtieth AAAI Conference on Artificial Intelligence* (AAAI Press, Washington, DC, USA, 2016).

[52] R. A. Fisher, On the mathematical foundations of theoretical statistics, *Philos. Trans. R. Soc. London. A* **222**, 309 (1922).

[53] M. J. Holland and K. Burnett, Interferometric detection of optical phase shifts at the Heisenberg limit, *Phys. Rev. Lett.* **71**, 1355 (1993).

[54] C. Gross, T. Zibold, E. Nicklas, J. Estève, and M. K. Oberthaler, Nonlinear atom interferometer surpasses classical precision limit, *Nature (London)* **464**, 1165 (2010).

[55] 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).

[56] R. Kaubruegger, P. Silvi, C. Kokail, R. van Bijnen, A. M. Rey, J. Ye, A. M. Kaufman, and P. Zoller, Variational spin-squeezing algorithms on programmable quantum sensors, *Phys. Rev. Lett.* **123**, 260505 (2019).