Methods

System Overview

Here we describe the experimental apparatus and some foundational techniques, with a system overview presented in Extended Data Fig. 1.

Loading. A cold ^{87}Rb atom cloud is loaded in a magneto-optical trap (MOT) inside a glass vacuum cell. Single atoms are then loaded stochastically into 852-nm optical tweezers generated by a high-speed spatial light modulator (SLM) and rearranged into a defect-free array using the SLM-based, AI-enabled parallel rearrangement protocol reported in our previous work[8]. Here, we choose a one-dimensional configuration of five well-isolated atom pairs, with an intra-pair spacing of 3 μm to ensure strong Rydberg interactions and an inter-pair separation of 30 μm to suppress crosstalk between different pairs. We then perform three-dimensional D2-line Lambda-enhanced gray molasses cooling[25] using three pairs of counter-propagating laser beams shared with the MOT beams, cooling the atoms to ~8 μK in 1-mK-deep traps.

Initialization. Physical qubits are encoded in the hyperfine clock states $|0\rangle = |F = 1, m_F = 0\rangle$ and $|1\rangle = |F = 2, m_F = 0\rangle$. Following the state preparation protocol in ref. 17, the initialization sequence begins with atoms optically pumped to the $|F = 2, m_F = 2\rangle$ stretched state in a bias magnetic field. Then, we rotate the magnetic field by 90 degrees and use the 795-nm Raman laser to drive $\sigma^{\pm}$ transitions from $|F = 2, m_F = 2\rangle$ to $|F = 1, m_F = 1\rangle$ and then to $|F = 2, m_F = 0\rangle$, using a Knill composite π pulse[41]. Finally, we rotate the magnetic field back to its original orientation, defining the quantization axis for the subsequent circuit. To further enhance the state preparation fidelity, we designate the $|F=1,m_F=0\rangle$ state as our target initial state, as this configuration naturally accommodates a highly efficient Raman-assisted pumping (RAP) scheme. We first apply a π pulse to transfer the atoms from $|F=2,m_F=0\rangle$ to $|F=1,m_F=0\rangle$. This is immediately followed by the RAP sequence, where we repeatedly apply π pulses on the $|F=1,m_F=\pm1\rangle\rightarrow|F=2,m_F=\pm1\rangle$ transitions, interleaved with depumping pulses resonant with the $|5S_{1/2},F=2\rangle\rightarrow|5P_{3/2},F'=2\rangle$ transition. The RAP protocol efficiently pumps any residual population in other ground-state sublevels into the target $|F=1,m_F=0\rangle$ state.

We note that the three-dimensional gray molasses cooling and multi-stage state preparation described above—out-of-circuit operations—are performed only before the circuit execution.

Single-qubit gates. The 795-nm Raman laser is blue-detuned by 200 GHz from the D1 line. Light is phase-modulated by an electro-optic modulator (EOM) driven with a 6.8-GHz IQ-modulated microwave signal, and a chirped Bragg grating converts this into amplitude modulation[27]. The Raman beam propagates along the quantization axis defined by the magnetic field and drives π transitions within the hyperfine clock manifold. We add an SLM to precisely align the Raman beam with the atom array, following the same alignment procedure as for the Rydberg beams described below.

We implement single-qubit X rotations using robust BB1 pulses[42] at a Raman Rabi frequency of 0.5 MHz, whereas Z rotations are realized by shifting the microwave phase. Arbitrary single-qubit rotations are decomposed into $Z - X - Z$ sequences. Through randomized benchmarking, we extract a single-qubit gate fidelity of 99.97%, which is primarily limited by Raman scattering and can be further improved by increasing the detuning.

Rydberg-mediated CZ gate

Rydberg excitation. Atoms in $|1\rangle$ are coupled to the $70S_{1/2}$ Rydberg state in a two-photon excitation scheme, using counter-propagating 420-nm σ^{-}-polarized and 1013-nm σ^{+}-polarized light via the $6P_{3/2}$ intermediate state. Through the intermediate state, atoms can be coupled to both $m_j = \pm 1/2$ Rydberg levels. We choose the $m_j = -1/2$ state for its stronger coupling, and apply a large magnetic field of 20 G to suppress unwanted coupling to the $m_j = +1/2$ state (Extended Data Figs. 1 and 4).

Our Rydberg lasers are built upon fiber seeds and fiber amplifiers, featuring excellent phase noise performance. The 420-nm laser delivers up to 4 W, generated from two amplified fiber seeds through frequency summing and doubling. The 1013-nm laser is driven by a fiber seed and amplified to a maximum output of 100 W. Both lasers are locked to an ultra-low-expansion cavity (finesse of 50,000 at 1013 nm and 840 nm).

In our experiment, we operate at a two-photon Rabi frequency $\Omega = 2\pi \times 5$ MHz under a large intermediate-state detuning of $\Delta = 2\pi \times 9.1$ GHz. Combined with an independent measurement of the 420-nm ground-state light shift—either by performing a Rydberg Ramsey experiment with the 420-nm laser switched off or by tracking the resonance shift versus 420-nm power—we estimate the single-photon Rabi frequencies to be $\Omega_{420} = 2\pi \times 492$ MHz and $\Omega_{1013} = 2\pi \times 185$ MHz. We note a pronounced imbalance between Ω_{420} and Ω_{1013}, which arises because the available 1013-nm power is limited by the stimulated Brillouin scattering threshold[43] in our 5-m-long delivery fiber. This imbalance introduces notably higher intermediate-state scattering (as discussed below), and can be straightforwardly addressed in the future by using a shorter delivery fiber.

Time-optimal CZ gate. The CZ gate is implemented by phase-modulating the 420-nm laser with a parameterized time-optimal pulse[17,28], $\phi(t) = A\cos(\omega t - \varphi_0) + \delta_0 t$. To physically realize the phase modulation, we employ a two-channel arbitrary waveform generator (AWG) to drive two tandem acousto-optic modulators (AOMs)—with a 4f lens system inserted between them—to imprint the sinusoidal phase profile onto the 420-nm beam. Each AOM is operated at a center frequency of +240 MHz, and the total phase modulation is distributed equally across the two AOMs, with a precisely adjusted relative delay to ensure synchronous modulation. This tandem configuration compensates for the diffraction angle shift inherent to phase/frequency modulation, thereby mitigating any unwanted amplitude modulation.

Intensity stabilization. Stable optical intensity is essential for achieving high-fidelity CZ gates. To this end, the 1013-nm laser is actively intensity-stabilized throughout the CZ gate execution. For the 420-nm path, the laser is first intensity-stabilized by a dedicated AOM and then sent through the two modulating AOMs described above. In each experimental cycle, we monitor the 420-nm power immediately before the atoms and use this measurement to update the intensity-stabilization setpoint for the next shot. This shot-to-shot update strategy compensates for slow drifts, ensuring reproducible optical power. We achieve a relative intensity stability of 0.2% (standard deviation over the mean) for both beams.

Rydberg beam alignment. We add an SLM for each Rydberg beam to precisely align the beam centers to the atom array (Extended Data Fig. 1a). Given that we currently use Gaussian beams for both paths (waist radii ~30 μm), if misaligned, atomic thermal motion would lead to significant decoherence from fluctuations in the Rabi frequencies and light shifts.

To address this, we develop a reproducible multi-axis spatial alignment protocol. By applying grating and Fresnel lens holograms to the SLM, we can control the beam position in three dimensions. We probe the atoms with fixed, long-duration pulses while scanning the beam position, which translates the spatially varying intensity into a multi-peak probability distribution. Specifically, for the 420-nm beam, we employ a Ramsey sequence to evaluate the differential light shift between states $|F=1, m_F=0\rangle$ and $|F=2, m_F=0\rangle$, which is proportional to the laser intensity (Extended Data Fig. 2a). For the 1013-nm beam, we employ a similar Ramsey sequence—incorporating an echo due to the much smaller differential light shift—along the radial direction of the tweezers. However, along the weakly confining axial direction of the tweezers, the alignment relies on a distinct mechanism: the strong 1013-nm ground-state light shift creates a deep attractive potential, pulling the atoms out of the tweezers if the beam center is misaligned (Extended Data Fig. 2b).

In practice, these spatial alignments are typically performed as independent one-dimensional sweeps. Beyond this, by executing a two-dimensional scan (transverse and longitudinal) and analyzing the signal for each individual atom, we can extract any angular tilt between the beam propagation axis and the one-dimensional atom chain. For routine daily maintenance, we perform automated transverse alignments for each beam to efficiently compensate for thermal and mechanical drifts. Notably, the SLMs used here also allow for future flat-top beam shaping.

Randomized benchmarking. We have characterized the CZ gate performance using four distinct randomized benchmarking (RB) configurations, which differ by whether an X-gate echo is inserted (to cancel the intrinsic single-qubit phase of the CZ gate), and by the choice of readout method (destructive pushout or loss-resolved measurement).

To calibrate the five gate parameters (A, ω, φ_0, δ_0, and duration T) describing the time-optimal pulse, we employ echoed RB sequences containing 32 entangling gates with destructive measurement, scanning the parameters and fitting the results to find the optimal values. Following a similar procedure, we calibrate the residual single-qubit phase for the non-echoed RB by scanning the global-Z compensation phase within the sequence (Extended Data Figs. 3a and 3b).

These optimized parameters are then fixed for all subsequent RB experiments, which are performed on five atom pairs by varying the number of CZ gates N (up to 64). The resulting decay curves for the echoed configuration are presented in Fig. 2e (destructive measurement) and Fig. 3d (loss-resolved measurement), while the corresponding non-echoed results are shown in Extended Data Fig. 3c. These data are fitted to an exponential decay of the form $ap^N + b$. For loss-resolved measurement, we fix $b = 1/4$ assuming a depolarizing channel, whereas for destructive measurement, we set $b = 0$. The gate fidelity is then extracted using the relation[20] $1 - F = (1 - p)(1 - b)$.

CZ gate simulation and error budget

Constructing the model. To analyze the error sources of the CZ gate and to obtain a breakdown of the error types, we simulate the full CZ gate dynamics using a complete density-matrix evolution that includes ten atomic levels, as shown in Extended Data Fig. 4. These levels comprise the two computational basis states $|0\rangle$ and $|1\rangle$, the auxiliary states $|0_L\rangle$ and $|1_L\rangle$ that represent leakage into other hyperfine Zeeman sublevels, three intermediate states, the target Rydberg state $|r_1\rangle$, an additional nearby Rydberg state $|r_2\rangle$ $(m_j = +1/2)$, as well as nearby Rydberg P states $|r_P\rangle$ (primarily 69P and 70P) that are populated via blackbody-radiation-induced transitions from the

Rydberg manifold. In the simulations, the single-photon Rabi frequencies, the single-photon detuning, the atomic spacing, and all other relevant parameters are chosen to match the experimental values. We explicitly consider the coupling between $|0\rangle$, $|1\rangle$ and the intermediate states, as well as the coupling between $|r_1\rangle$, $|r_2\rangle$ and the intermediate states.

Error sources and breakdown. To enable a quantitative analysis of the relative contributions of Pauli errors and leakage errors, we carefully account for the branching ratios associated with radiative decay from both the intermediate and Rydberg states, incorporating nearly all dominant two-photon and four-photon emission processes (the combined branching ratio for higher-order processes is estimated to be below 0.005). We assume a min-to-max rise/fall time of approximately 20 ns in a Blackman profile, which plays a critical role in the population dynamics of the intermediate states and the nearby Rydberg state $|r_2\rangle$. We introduce Gaussian-distributed detuning noise in the simulations to model the experimentally measured single-qubit ground-Rydberg dephasing time T_2^*. This effective T_2^* captures the combined effects of Doppler shifts arising from the finite atomic temperature, light-shift fluctuations induced by intensity fluctuations of the 420-nm and 1013-nm lasers, as well as electric-field fluctuations and laser phase noise. In contrast, blockade-strength fluctuations originating from finite-temperature-induced position spread of the atoms are treated separately, as they are not encompassed by the measured T_2^*. We additionally include the degradation of gate fidelity due to experimental parameter miscalibration.

Result analysis and outlook. The simulated results, summarized in Extended Data Table 1, are in good agreement with the experimentally observed CZ gate performance. A substantial fraction of the total error during the CZ gate arises from atom-loss events, which can be converted into erasure errors through direct loss detection, thereby effectively improving the fidelity. Moreover, in the current experimental configuration, the strong imbalance between the single-photon Rabi frequencies at 420 nm and 1013 nm leads to significantly higher scattering from the intermediate states. By adopting a more balanced configuration, we expect an immediate improvement of ~0.1% in the raw fidelity. Looking forward, further increasing the Rabi frequency and detuning, extending T_2^*, and applying a stronger magnetic field to suppress coupling to the other m_j Rydberg state are expected to enable gate fidelities approaching 99.9% in the near future.

Mid-circuit non-destructive measurement

Closed-transition imaging. As the primary mechanism of the closed-transition imaging is introduced in the main text, here we detail the strobing sequence. Within a single modulation cycle, the tweezers are temporarily switched off; the counter-propagating imaging beams are then pulsed sequentially (not simultaneously), after which the tweezers are immediately switched back on. This MHz-rate rapid switching is controlled by AWGs driving AOMs, into which the laser beams are focused. To minimize atom loss probability and maximize state-resolved fidelity, we systematically optimize the strobing period, duty cycle, laser intensities, trap depth, and total repetition number, with the optimal parameters summarized in Extended Data Table 2.

Mid-circuit PGC imaging. After closed-transition imaging, we implement finite-field one-dimensional local polarization gradient cooling (local PGC) to simultaneously cool and detect the retained atoms. Our setup consists of a pair of counter-propagating $\sigma^+ + \sigma^-$ beams, with a weak repump beam combined into the σ^+ beam path. We note that the local PGC beams are also used for local gray molasses cooling and for depumping in Raman-assisted pumping (discussed below), with different operating frequencies.

To operate under a finite magnetic field, following the scheme in ref. 13, we introduce a relative frequency detuning between the two cooling beams. In a reference frame where the polarization vector rotates in time, the relative detuning generates a fictitious magnetic field that cancels the external bias field, thereby restoring the PGC mechanism. Alternatively, viewed from the individual Zeeman sublevels, the σ^+ and σ^- beams respectively address the lower and higher energy states separated by $\Delta m_F = 2$, with their frequency difference precisely matching the $2\mu_B g_F B_{ext}$ Zeeman gap. This relative detuning is also applied to local gray molasses cooling. In particular, the local PGC lasers are intentionally red-detuned by ~58 MHz from the $F = 2 \rightarrow F' = 3$ transition to improve cooling performance.

AI model for atom recognition. To improve imaging fidelity and generalize to various imaging configurations, we develop an atom recognition model based on a convolutional neural network. By leveraging spatial correlations across neighboring pixels, the model can accurately determine atomic occupancy even with low fluorescence photon counts, thereby allowing a significantly shorter imaging time. The model features a lightweight, five-layer fully convolutional architecture with approximately 20k trainable parameters.

To optimize class separability, we employ a Fisher ratio loss function, which is designed to maximize the inter-class distance between atomic presence and vacancy while simultaneously minimizing intra-class variance. The loss is defined as:

$$L_{Fisher} = -\frac{(\mu_1 - \mu_0)^2}{\sigma_1^2 + \sigma_0^2}$$

where μ and σ^2 represent the mean and variance of the two classes.

To ensure the model's versatility, training is performed across three distinct experimental datasets, each containing over 1000 images. We summarize the corresponding imaging configurations and the procedures used to obtain ground-truth labels as follows: (1) Global imaging: This refers to conventional PGC imaging operated at zero magnetic field. Owing to the well-separated bimodal fluorescence histogram, labels are assigned by simple thresholding. (2) Closed-transition imaging: Atoms are deterministically prepared in $|0\rangle$ or $|1\rangle$, followed by closed-transition imaging and a subsequent global imaging shot. The known prepared state together with the post-check is used to assign the ground-truth label. (3) Local PGC imaging: A global imaging shot is taken immediately before the local PGC imaging to determine atom occupancy, and this pre-check is used as the label.

We illustrate the raw photon count histograms alongside the AI-processed intensity distributions for all three imaging configurations in Extended Data Fig. 5. The AI model consistently exhibits excellent class separability. Given this, we estimate the imaging time for each configuration can be substantially reduced by a factor of two to three.

State-resolved fidelity and survival probability. To characterize the mid-circuit non-destructive measurement, we prepare atoms in $|0\rangle$ or $|1\rangle$ and perform the two-stage imaging. The joint outcomes of the two imaging stages are shown as the correlation scatter plot in Fig. 3c, which is used to resolve $|0\rangle$, $|1\rangle$, and loss events. However, these results are affected by state preparation errors. To correct for this, we perform an additional calibration in which atoms prepared in $|0\rangle$ or $|1\rangle$ are measured with and without a pushout pulse. Combining the state preparation errors—extracted from the above four configurations—with the results in Fig. 3c, we perform a maximum-likelihood analysis to obtain the intrinsic imaging metrics. We derive a state-resolved fidelity of $99.02 \pm 0.10\%$ $(98.68 \pm 0.11\%)$ and a survival probability of $99.73 \pm 0.07\%$ $(99.67 \pm 0.06\%)$ for the dark (bright) state. We expect these values could be further improved by reducing the imaging time, enabled by the excellent performance of AI recognition.

Mid-circuit cooling

Local GM. Although local PGC provides effective imaging and preliminary cooling, we develop a more efficient local gray molasses (local GM) scheme. Generally, GM cooling operates on a Λ-type three-level system, wherein two coherent fields couple two ground states to an excited state with a blue detuning[44]. This creates a dark state and a spatially varying bright state, driving a Sisyphus-like cooling cycle. While our 3D global GM utilizes standard cooling and repump beams to couple two hyperfine states, here, operating under a finite magnetic field, we employ two cooling beams with a relative detuning to coherently couple two Zeeman sublevels.

Experimentally, this is realized using the same counter-propagating $\sigma^+ + \sigma^-$ beams (also with a weak repump component) as for local PGC, but blue-detuned by ~60 MHz from the $F = 2 \rightarrow F' = 2$ transition. After optimizing the laser frequency, relative detuning, and beam intensities, we measure a local GM temperature of ~12 μK via release–recapture, close to the ~8 μK of our global GM (see Extended Data Fig. 6).

We initially anticipated that this local GM would be sufficient to sustain the CZ gate fidelity throughout multiple mid-circuit rounds. However, as shown in Fig. 4c, a notable degradation persists. While the exact mechanism warrants further investigation, we hypothesize that this performance drop is primarily driven by the uncooled dimensions, particularly the axial direction.

Mid-circuit RSC. To provide a more efficient entropy removal mechanism, we develop a mid-circuit Raman sideband cooling (RSC) protocol. As illustrated in Fig. 4a, the cooling is driven by three beams derived from our 795 nm single-qubit Raman laser system (200 GHz blue-detuned from the D1 line), with RSC1 phase-modulated by an EOM, while RSC2 and RSC3 remain unmodulated. To isolate the desired Raman transition from the multiple sidebands of RSC1, we introduce a frequency offset by setting RSC2 and RSC3 30 MHz above the RSC1 carrier. In this configuration, only the +1 sideband of RSC1 participates in the Raman transition (coupling $|1,0\rangle$), acting in concert with RSC2 and RSC3 (coupling $|2,0\rangle$). Specifically, RSC1 and RSC3 address the radial motional mode, while RSC1 and RSC2 address both the radial and axial motional modes.

The cooling cycle begins with a Raman sideband transfer from $|1,0\rangle$ to $|2,0\rangle$, which decreases the phonon number. For this step, we employ a Blackman pulse profile[45] to suppress off-resonant transitions. Given the broad initial phonon distribution—due to a lack of effective cooling after mid-circuit measurements—we apply a dynamically chirped pulse sequence to enhance the capture range of our RSC. This sequence initiates with high-order sideband transitions and short transfer pulses to couple states with high phonon numbers, then gradually sweeps toward first-order transitions and longer pulses for low-phonon states. Following the Raman transfer step, the cooling cycle is closed by Raman-assisted pumping (RAP, as detailed in the earlier Methods section). The RAP dissipates the motional energy and resets the atoms back to $|1,0\rangle$, while preserving the lowered phonon number. Notably, the depumping process in RAP reuses the local PGC beams ($\sigma^{+}+\sigma^{-}$), with the frequency tuned to the $F=2\rightarrow F'=2$ transition. We emphasize that our RSC protocol simultaneously achieves mid-circuit cooling and state initialization.

References

41. Souza, A. M., Álvarez, G. A. & Suter, D. Robust Dynamical Decoupling for Quantum Computing and Quantum Memory. Phys. Rev. Lett. 106, 240501 (2011).

42. Vandersypen, L. M. K. & Chuang, I. L. NMR techniques for quantum control and computation. Rev. Mod. Phys. 76, 1037–1069 (2005).

43. Vidal, S. et al. 20 µm-core polarization maintaining endlessly single mode photonic crystal fiber for delivery of high-power single frequency lasers. Opt. Express 33, 15603 (2025).

44. Weidemüller, M., Esslinger, T., Ol'shanii, M. A., Hemmerich, A. & Hänsch, T. W. A Novel Scheme for Efficient Cooling below the Photon Recoil Limit. Europhys. Lett. EPL 27, 109–114 (1994).

45. Kasevich, M. & Chu, S. Laser cooling below a photon recoil with three-level atoms. Phys. Rev. Lett. 69, 1741–1744 (1992).

Extended Data Figures and Tables

Extended Data Fig. 1 | Experimental setup and control architecture. a, Schematic of the optical layout. Single ^{87}Rb atoms are trapped in 852-nm optical tweezers and deterministically arranged using an SLM-based, AI-enabled parallel rearrangement protocol. The diagram illustrates the beam delivery and polarization geometries for Rydberg gates, Raman single-qubit rotations, Raman sideband cooling (RSC), closed-transition imaging, local polarization gradient cooling (local PGC) and local gray molasses cooling (local GM). SLMs are employed in the Raman and Rydberg paths for precise multi-axis beam alignment. **b,** Relevant ^{87}Rb energy levels and transitions driven by the optical setup. **c,** Control electronics. The architecture comprises four two-channel AWGs referenced to a common external clock, with relative delays precisely calibrated across all channels to ensure strict synchronization. For time-optimal Rydberg gate, the Rydberg AWG imprints a sinusoidal phase profile onto the 420-nm laser; this modulation is split equally across two channels driving a tandem AOM configuration. Additional AWGs independently drive the 6.8-GHz microwave IQ modulation, the fast switching of the tweezer and Raman beams, and the strobing of the 780-nm closed-transition imaging laser. (*) The RSC2 beam propagates perpendicular to the array plane and is focused along the vertical direction to spatially confine the operation region. The beam has vertical linear polarization that, relative to the quantization axis, decomposes into σ^+ and σ^- components, with the former participating in the Raman sideband cooling process.

Extended Data Fig. 2 | Rydberg beam alignment. Three-dimensional reproducible positioning of the beam centers is achieved by combining holographic gratings and Fresnel lens phase patterns on the SLMs. **a,** Spatial alignment via Ramsey measurement. The state $|0\rangle$ population is measured after a fixed free-evolution time (set to an integer multiple of the oscillation period) while sweeping the beam position. The experimental data (blue points) are fitted to a model (solid orange line) that accounts for the intensity-dependent Ramsey phase accumulation rate across the Gaussian beam profile. **b,** Spatial alignment via 1013-nm-induced atom loss, along the axial direction of the tweezers. The strong 1013-nm ground-state light shift creates a deep attractive potential, overcoming the weak axial confinement and pulling the atoms out of the tweezers, if the beam center is misaligned. The exact beam position is determined by sweeping the beam and measuring the survival fraction (blue points). A truncated Gaussian fit (solid orange line) precisely pinpoints the beam center.

Extended Data Fig. 3 | Non-echoed randomized benchmarking of the CZ gate. a, Circuit used to benchmark the non-echoed CZ gate. Interleaved Z-rotations, denoted as $Z(\phi)$, are incorporated to compensate for the single-qubit phase accumulated during the time-optimal Rydberg pulses. The final recovery unitary, U_{rec}, comprises five single-qubit rotations R_{fi} $(i = 1, \dots, 5)$ separated by four CZ gates. The parameters of R_{fi} are determined by a least-squares numerical optimization to invert the preceding circuit, thereby mapping the overall sequence to the identity operation. **b**, Example calibration of the compensation phase. The single-qubit phase ϕ is scanned to maximize the return probability $P_{|00\rangle}$ after the application of 32 entangling gates. **c**, Benchmarking results averaged over five atom pairs. Blue data points represent the raw experimental results obtained via destructive pushout measurement, yielding a fidelity of 99.58(1)%. Orange data points show the loss-corrected results acquired using non-destructive measurement, which enables post-selection against atom loss events to reach a fidelity of 99.80(4)%. Solid lines denote exponential fits to the decay data, with the asymptote fixed to 1/4 for the loss-corrected data and 0 for the raw data.

Extended Data Fig. 4 | Ten-level atomic model for the CZ gate simulation. Schematic of the relevant energy levels and decay channels incorporated into the density-matrix evolution. The system comprises the two computational basis states $|0\rangle$

and $|1\rangle$, the auxiliary states $|0_L\rangle$ and $|1_L\rangle$ that represent leakage into other hyperfine Zeeman sublevels, three intermediate states, the target Rydberg state $|r_1\rangle$, an additional nearby Rydberg state $|r_2\rangle$ $(m_j = +1/2)$, as well as nearby Rydberg P states $|r_P\rangle$ (primarily 69P and 70P). The intermediate states have a lifetime of 110 ns, with spontaneous emission branching ratios given by $\eta_{e\to0} = 0.1698, \eta_{e\to1} = 0.1737, \eta_{e\to0_L} = 0.2483, \eta_{e\to1_L} = 0.4082$. For the Rydberg states, the total effective lifetime is 151.55 μs, derived from a combination of BBR decay (240.28 μs) and spontaneous radiative decay (410.41 μs). We assume the BBR decay components results purely in atom loss and radiative decay has branch ratios: $\eta_{r_1\to0} = 0.1960, \eta_{r_1\to1} = 0.2137, \eta_{r_1\to0_L} = 0.2256, \eta_{r_1\to1_L} = 0.3647$.

Extended Data Fig. 5 | AI-based atom recognition across different imaging configurations. a–**c**, Raw fluorescence histograms captured under global imaging, local PGC imaging, and closed-transition (CT) imaging. **d**–**f**, Corresponding AI-processed intensity distributions for global, local PGC, and CT imaging. Across all configurations, the AI model consistently exhibits excellent class separability.

Extended Data Fig. 6 | Characterization of gray molasses (GM) cooling performance. a, Optimization of local GM (1D; finite magnetic field) cooling: survival probability as a function of single-photon detuning and relative detuning, obtained through a 15-μs release-and-recapture measurement. The blue star identifies the optimal detuning parameters selected in this work. **b**, **c**, Release-and-recapture survival curves for global GM (3D; zero magnetic field) and local GM cooling. Solid lines represent Monte Carlo simulations, yielding fitted atomic temperatures of 8 μK and 12 μK for the global and local cases, respectively.

Extended Data Table 1 | Error budget of the two-qubit CZ gate. Detailed decomposition of error sources and error types, derived from density-matrix simulations under experimental conditions. The first two columns identify the specific physical mechanisms and quantify their respective infidelities. These errors are categorized into three distinct types: Pauli errors, leakage outside the computational basis, and physical atom loss, with their relative proportions detailed in the final three columns. The bottom "Total fidelity" row summarizes the cumulative gate performance, alongside the relative contribution of each error category.

Extended Data Table 2 | Optimized experimental parameters for closed-transition imaging. These values were calibrated to maximize the state-resolved fidelity while minimizing atom loss. The strobing duty cycle defines the active temporal fractions within a strobing period: 21% for the imaging beams and 79% for the optical tweezers. During the trap-off interval, the two counter-propagating imaging beams are toggled alternately with equal durations.

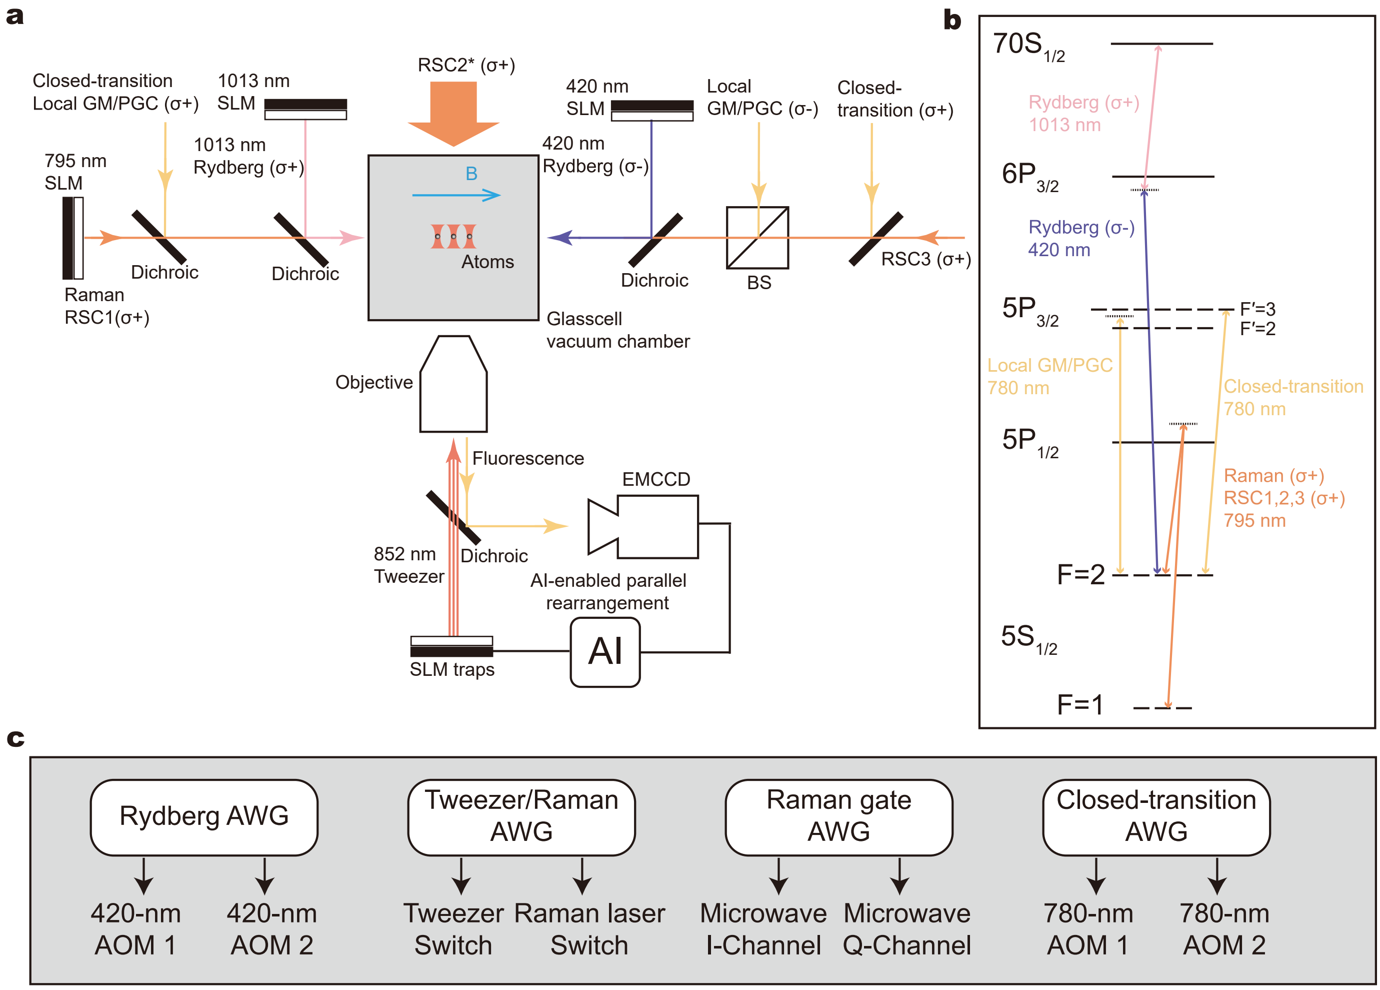

Extended Data Fig. 1

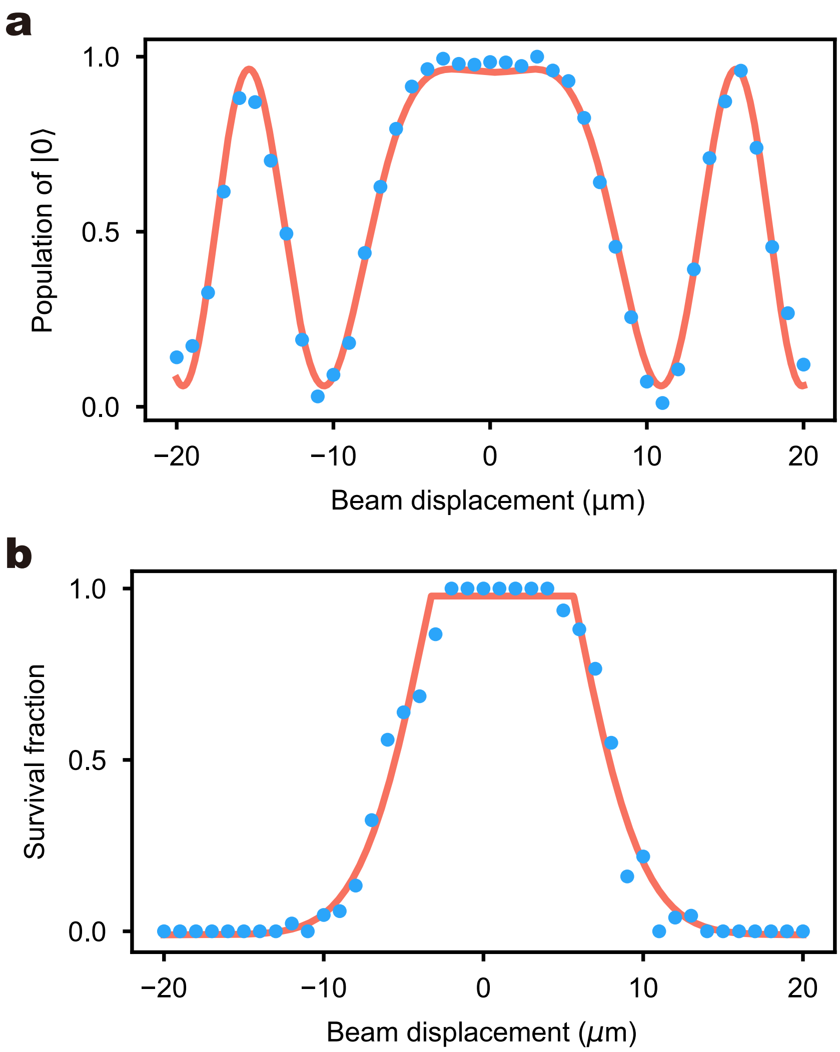

Extended Data Fig. 2

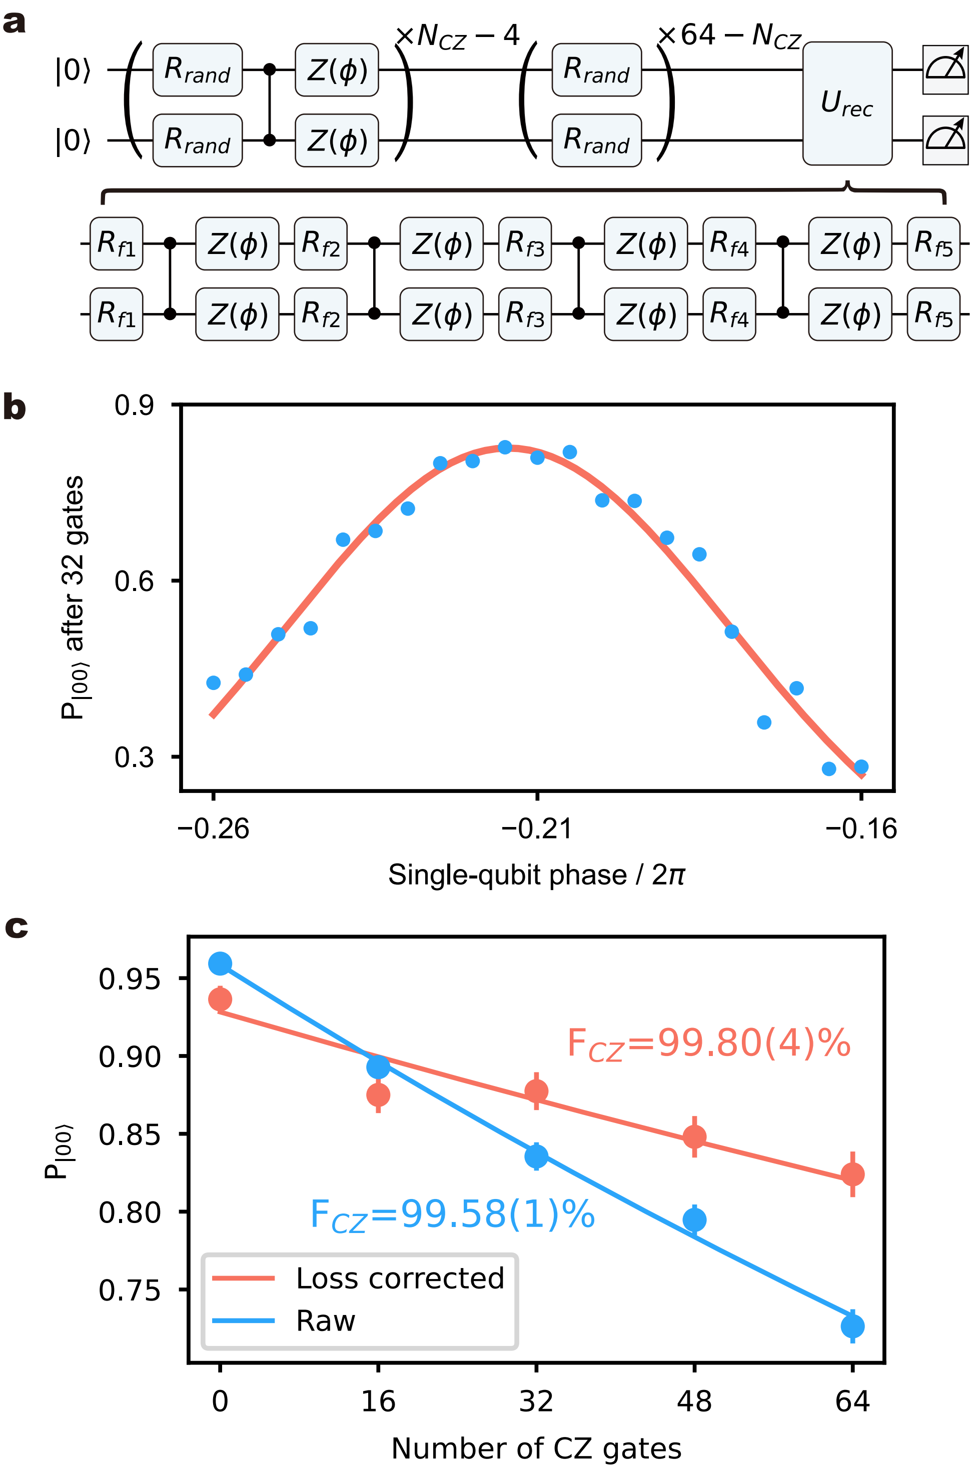

Extended Data Fig. 3

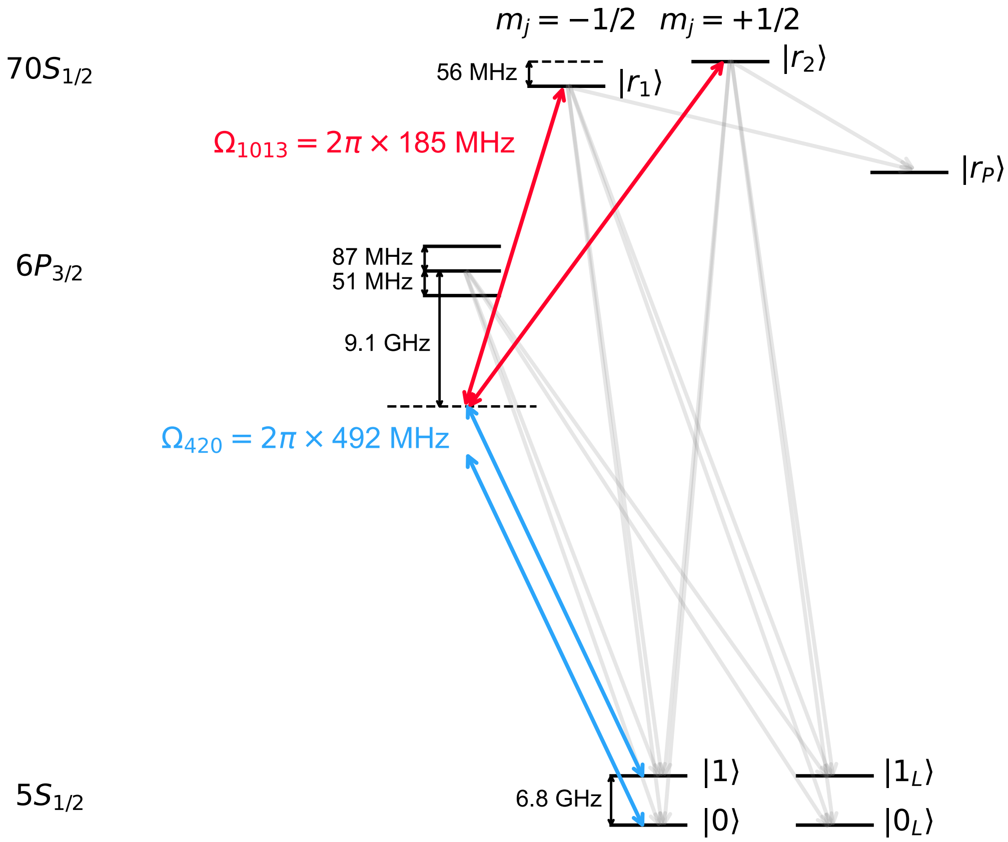

Extended Data Fig. 4

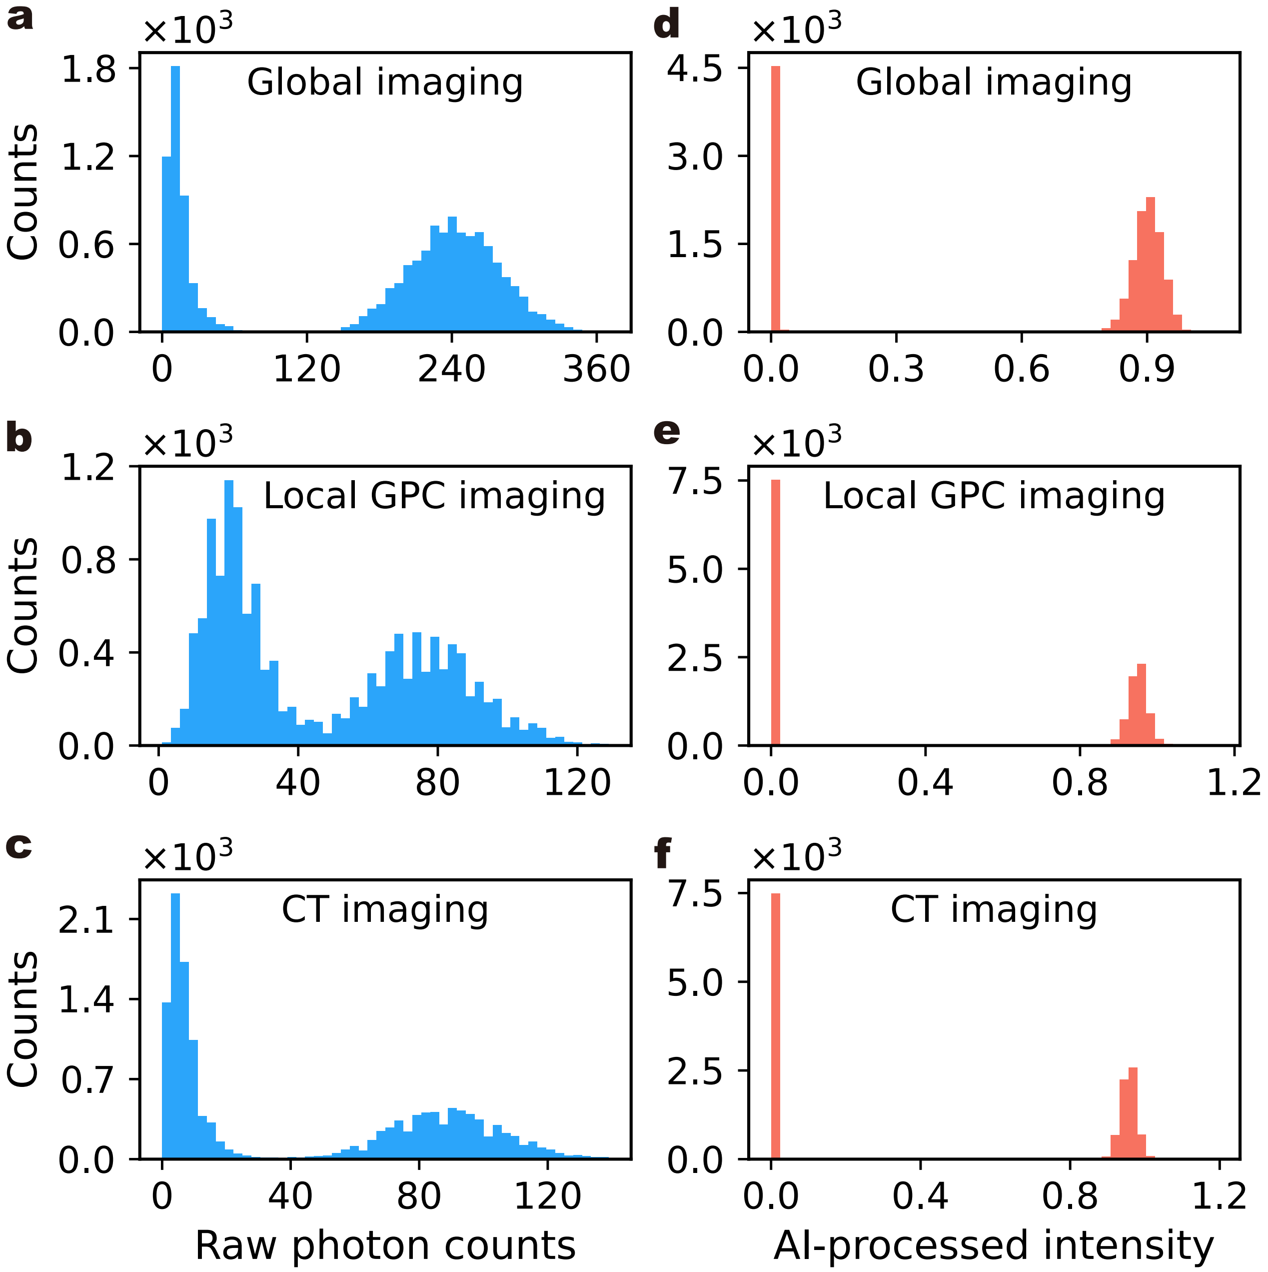

Extended Data Fig. 5

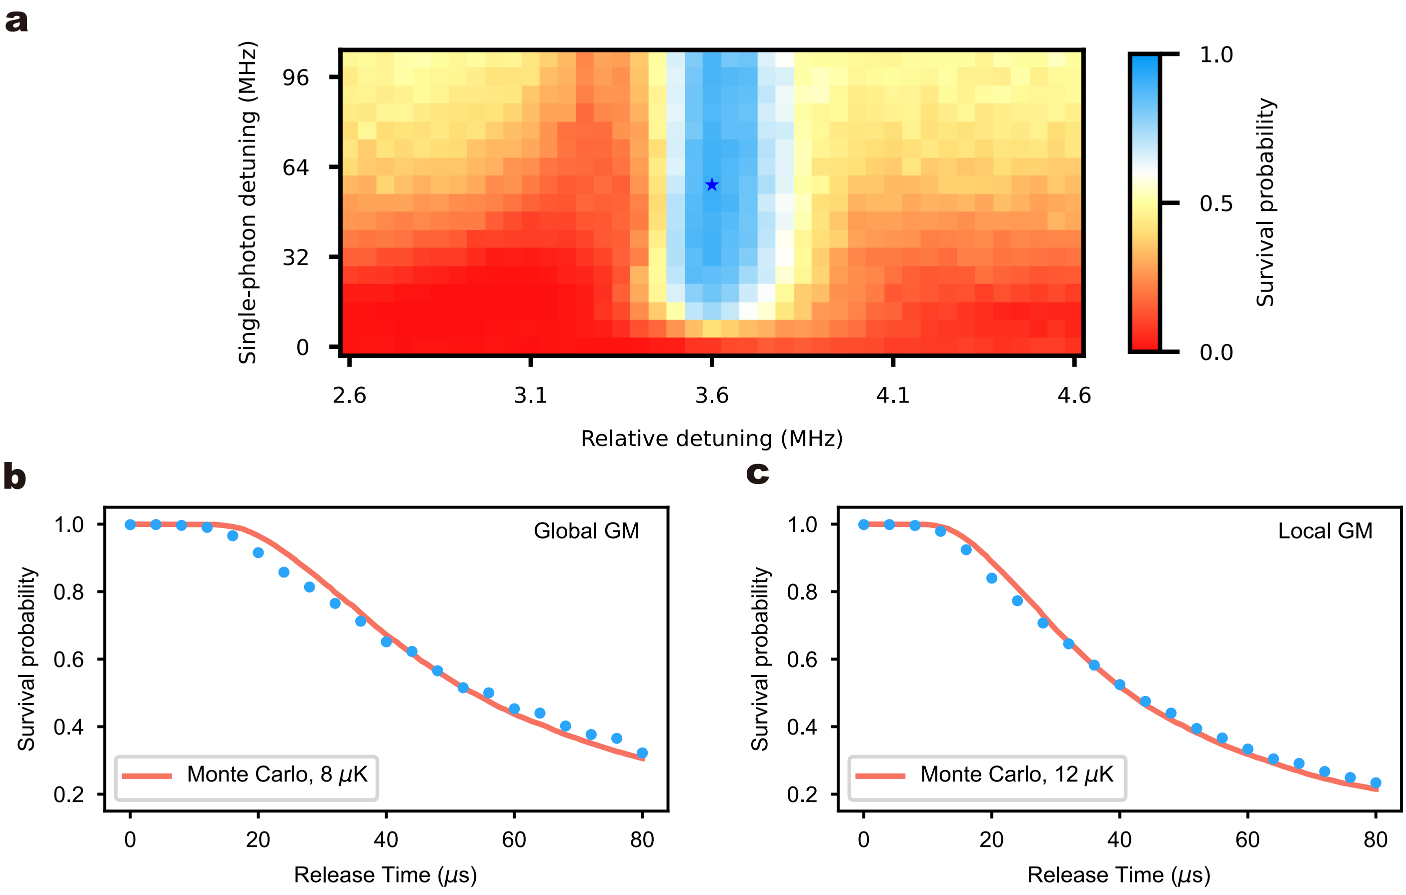

Extended Data Fig. 6

Error source		Pauli	Leakage	Loss
Scattering $\|1\rangle$	0.108%	20%	64%	16%
Scattering $\|0\rangle$	0.060%	17%	76%	7%
Rydberg $T_1 = 151.55\ \mu s$	0.054%	5%	15%	80%
Rydberg $T_2^* = 2\ \mu s$	0.082%	54%	0%	46%
Rydberg $m_j = +1/2$	0.004%	0%	2%	98%
Position fluct.	<0.001%	-	-	-
Params mis-calibration	0.092%	10%	1%	89%
Total fidelity	**99.600%**	**22%**	**31%**	**47%**

Extended Data Table. 1

Experimental parameter	Value
Trap depth	2.5 mK
Strobing period	0.89 µs
Strobing duty cycle	21% (79%) for imaging(trap)
Magnetic field	~2.5 G
Beam intensity	~1.2 I_{sat} per beam
Repetition	15,000

Extended Data Table. 2