Estimation of a general time-dependent Hamiltonian for a single qubit

L.E. de Clercq[1], R. Oswald[1], C. Flühmann[1], B. Keitch[1,†], D. Kienzler[1], H.-Y. Lo[1], M. Marinelli[1], D. Nadlinger[1], V. Negnevitsky[1] & J.P. Home[1]

The Hamiltonian of a closed quantum system governs its complete time evolution. While Hamiltonians with time-variation in a single basis can be recovered using a variety of methods, for more general Hamiltonians the presence of non-commuting terms complicates the reconstruction. Here using a single trapped ion, we propose and experimentally demonstrate a method for estimating a time-dependent Hamiltonian of a single qubit. We measure the time evolution of the qubit in a fixed basis as a function of a time-independent offset term added to the Hamiltonian. The initially unknown Hamiltonian arises from transporting an ion through a static laser beam. Hamiltonian estimation allows us to estimate the spatial beam intensity profile and the ion velocity as a function of time. The estimation technique is general enough that it can be applied to other quantum systems, aiding the pursuit of high-operational fidelities in quantum control.

[1] Institute for Quantum Electronics, ETH Zürich, Otto-Stern-Weg 1, 8093 Zürich, Switzerland. † Present address: Department of Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK. Correspondence and requests for materials should be addressed to J.P.H. (email: jhome@phys.ethz.ch).

Estimation of the underlying dynamics, which drive the evolution of systems is a key problem in many areas of physics and engineering. This knowledge allows control inputs to be designed, which account for imperfections in the physical implementation. For closed quantum systems, the time dependence of a system is driven by the Hamiltonian through Schrödinger's equation. If the Hamiltonian is static in time, a wide range of techniques have been proposed for recovering the Hamiltonian[1–4], which have been applied to a variety of systems including chemical processes[5] and quantum dots[6,7]. These methods often involve estimation of the eigenvectors and eigenvalues of the Hamiltonian via spectroscopy, or through pulse–probe techniques for which a Fourier transform of the time evolution gives information about the spectrum.

These methods are not directly applicable to time-dependent Hamiltonians, which are becoming increasingly important as quantum engineering pursues a combination of high-operational fidelities and speeds, often involving fast variation of control fields, which are particularly susceptible to distortion before reaching the quantum device[8–12]. The time-varying case has thus far been studied in cases where the variation is along a single dimension in the Hilbert space, which for the commonly studied spin is a single spatial direction. In the case that the measured fields dominate the evolution (strong field limit), measurement of the system evolution as a function of time suffices for the reconstruction. For fields which are weaker than other available control fields (weak-field limit) the latter can be used to modulate the effect of the signal Hamiltonian on the quantum system[13–15], providing an excellent signal-to-noise ratio. A further complication arises when a time-varying Hamiltonian contains non-commuting terms (for example, time-variation along two spin axes), because the evolution of the quantum system depends not only on their separate influences, but also on products arising from the non-commutativity. For unspecified time-dependent coefficients, no analytical solution to Schrödinger's equation exists[16,17]. In the weak-field limit, strong control fields can be used to separate out the different components using modulation, however, when the Hamiltonian itself is strong (as is the case in fast quantum control) these techniques cannot be applied.

In this article, we propose and demonstrate a method for reconstructing a general time-dependent single qubit Hamiltonian with non-commuting terms. The technique involves observing the evolution of the spin projection on the z-axis, while applying a static offset to one of the terms of the Hamiltonian. By varying the static offset, we build up data sets, which contain sufficient information to extract the full time-dependent Hamiltonian. Parameterizing the two time-dependent terms using basis splines (B-splines), we introduce an iterative fitting technique, which finds the Hamiltonian that best matches the data. We benchmark the reconstruction method experimentally by transporting a single trapped ion through a static laser beam, a technique suited to scaling up trapped-ion quantum information processing[18,19]. We perform two consistency checks on the Hamiltonian estimation using four separate reconstructions. For the first two, we compare two cases which use the same ion velocity profile, but different laser beam positions. For the second consistency check, we use the same laser beam, but change the velocity profile between the two by using different sets of time-varying control potentials. The method produces consistent experimental parameters in both cases, indicating the success of the reconstruction technique. Our method is applicable to spin Hamiltonians of the general form $\hat{H} = \sum_i f_i(t)\hat{\sigma}_i$, where the $f_i(t)$ are arbitrary time-dependent functions and $\hat{\sigma}_i$ are the Pauli operators.

Results

Hamiltonian estimation method. In our experiments, a Hamiltonian with two non-commuting time-dependent terms arises when we perform quantum logic gates by transporting an ion through a static laser beam[18,19]. In this case, the Hamiltonian describing the interaction between the ion and the laser can be written in an appropriate rotating frame as

$$\hat{H}_I(t) = \frac{\hbar}{2}\left(-\Omega(t)\hat{\sigma}_x + \delta(t)\hat{\sigma}_z\right) \tag{1}$$

which includes a time-varying Rabi frequency $\Omega(t)$, and an effective detuning $\delta(t)$, which is related to the first-order Doppler shift of the laser in the rest frame of the moving ion (see Methods for details).

To reconstruct the Hamiltonian, we make use of two additional capabilities. First, we can switch-off the Hamiltonian at time t_{off} on a timescale much faster than the qubit evolution. Subsequently measuring in the $\hat{\sigma}_z$ basis, we can obtain $\langle\hat{\sigma}_z(t_{off})\rangle$. On its own, this does not allow us to separate the contributions from $\Omega(t)$ and $\delta(t)$. To do this, we use a second capability, which is the ability to add a controlled offset $\hat{H}_s = \hbar\,\delta_L\hat{\sigma}_z/2$ to the Hamiltonian, resulting in $\hat{H}_I(t) + \hat{H}_s$. The resulting spin measurement is now dependent on both t_{off} and the set value of δ_L. Repeating the experiment for a range of values of δ_L but with otherwise identical settings, we obtain an estimate of the expectation value which we denote as $\langle\hat{\sigma}_z^{meas}(t_{off}, \delta_L)\rangle$.

Hamiltonian extraction involves finding the functions $\delta(t)$ and $\Omega(t)$, which generate spin populations $\langle\hat{\sigma}_z^{sim}(t_{off}, \delta_L)\rangle$ that most closely match the data. We minimize the reduced χ^2 cost function.

$$J = \frac{1}{v}\sum_{t_{off}}\sum_{\delta_L}\left[\frac{\langle\hat{\sigma}_z^{meas}(t_{off}, \delta_L)\rangle - \langle\hat{\sigma}_z^{sim}(t_{off}, \delta_L)\rangle}{\sigma^{meas}(t_{off}, \delta_L)}\right]^2 \tag{2}$$

where $v = N - n - 1$ is the number of degrees of freedom, with N the number of data points, n the number of fitting parameters and $\sigma^{meas}(t_{off}, \delta_L)$ the s.e. on the estimated $\langle\hat{\sigma}_z^{meas}(t_{off}, \delta_L)\rangle$. This is subject to the initial condition $|\Psi(t=0, \delta_L)\rangle = |0\rangle$, and the following restrictions, which are imposed by quantum mechanics

$$\begin{aligned}i\hbar\frac{\partial}{\partial t}|\Psi(t, \delta_L)\rangle &= (\hat{H}_I(t) + \hat{H}_s)|\Psi(t, \delta_L)\rangle, \\ \langle\hat{\sigma}_z^{sim}(t, \delta_L)\rangle &= \langle\Psi(t, \delta_L)|\hat{\sigma}_z|\Psi(t, \delta_L)\rangle\end{aligned} \tag{3}$$

for all δ_L.

One challenge in obtaining an estimate for the Hamiltonian is that we must optimize over continuous functions $\delta(t)$ and $\Omega(t)$. To address this, we represent $\delta(t)$ and $\Omega(t)$ with a linear combination of B-spline polynomials, which allow the construction of smooth functions using only a few parameters[20]. Any smooth function $S(t)$ can be written in terms of B-spline polynomials $B_{i,k}(t)$ and a set of weights α_i as

$$S(t) = \sum_{i=0}^{n}\alpha_i B_{i,k}(t). \tag{4}$$

The polynomial B-spline functions $B_{i,k}(t)$ are of order k with each polynomial centred at a time t_i, which is parameterized by the index i. Further details and an example can be found in Methods. Using the B-spline form for $\delta(t)$ and $\Omega(t)$, the cost function is minimized by adjusting the weights of the B-spline decomposition. Solving this optimization problem in general is hard, because it is non-linear and non-convex due to the nature of Schrödinger's equation and the use of projective measurements. This produces a non-trivial relation between the weights and the spin populations as discussed in Methods. To overcome this challenge, we have implemented a method which we call 'Extending the Horizon Estimation' (EHE) in analogy to a

well-established technique called 'Moving Horizon Estimation' (MHE)[21].

The key idea is that because our measurement data arises from a causal evolution, we can also estimate the Hamiltonian in a causal way. Instead of optimizing J over the complete time span at once, we first restrict ourselves to a small, initial time span reaching only up to the start of the qubit dynamics $0 < t_{off} < T_0$, where we denote T_0 as the time horizon for the first step. Optimizing J over this short time span requires fewer optimization parameters and is simpler than attempting to optimize over the full data set. Once we have solved this small sub-problem, we extend the time horizon to T_1 where $T_1 = T_0 + \tau$ and re-run the optimization, extrapolating the results of the initial time span into the extended region to provide good starting conditions for the subsequent optimization. This procedure is iterated until the time span extends over the whole data set ($T_{n_{max}} = \max(t_{off})$). The method allows us to reduce the number of B-spline functions used to represent $\delta(t)$ and $\Omega(t)$, and also reduces the amount of data considered in the early stages of the fit, when the least is known about the parameters. This facilitates the use of non-linear minimization routines, which are based on local linearization of the problem and converge faster near the optimum. More details regarding the optimization routine can be found in Methods.

Conceptually EHE is very similar to MHE. The main difference is that in MHE the time span has a fixed length and thus its origin is shifted forward in time along with the horizon. In EHE, the origin stays fixed at the expense of having to increase the time span under consideration. MHE avoids this by introducing a so-called arrival cost to approximate the previous costs incurred before the start of the time span. This keeps the computational burden fixed over time, which is very important as MHE is usually used to estimate the state of a system in real-time, often on severely constrained embedded platforms. Since neither constraint applies to our problem, we decided to extend the horizon rather than finding an approximate arrival cost. This is advantageous since finding the arrival cost in the general case is still an open problem[22]. Due to the similarity between MHE and EHE, we anticipate future improvements by adapting techniques used in MHE to EHE. This might be used to reduce the data-processing required for reconstruction, which for EHE scales as $T_{n_{max}}^2$.

Experimental implementation. To test the ability of the method to reliably extract a Hamiltonian from data, we apply it to the Hamiltonian for an trapped-ion qubit during transport through a near-resonant laser beam. Our qubit is encoded in the electronic states of a calcium ion, which is defined by $|0\rangle \equiv \left|{}^2S_{1/2}, M_J = 1/2\right\rangle$ and $|1\rangle \equiv \left|{}^2D_{5/2}, M_J = 3/2\right\rangle$. This transition is well-resolved from all other transitions, and has an optical frequency of $\omega_0/(2\pi) \simeq 411.0420$ THz. The laser beam points at $45°$ to the transport axis, and has an approximately Gaussian spatial intensity distribution. The time-dependent velocity $\dot{z}(t)$ of the ion is controlled by adiabatic translation of the potential well in which the ion is trapped. This is implemented by applying time-varying potentials to multiple electrodes of a segmented ion trap, which are generated using a multi-channel arbitrary-waveform generator, each output of which is connected to a pair of electrodes via a passive third-order low-pass Butterworth filter. The result is that the ion experiences a time-varying Rabi frequency $\Omega(t)$ and a laser phase which varies with time as $\Phi(t) = \phi(z(t)) - \omega_L t$, where $\phi(z(t)) = k_z(z(t))z(t)$ with $k_z(z(t))$ the laser wave vector projected onto the transport axis at position $z(t)$ and ω_L the laser frequency. The spatial variation of $k_z(z(t))$ accounts for the curvature of the wavefronts of the Gaussian laser beam. To create a Hamiltonian of the form of equation (1), we

work with the differential of the phase, which gives a detuning $\delta(t) = \delta_L - \dot{\phi} = \delta_L - (k'_z(z)z + k_z(z))\dot{z}$ with $\delta_L = \omega_L - \omega_0$ the laser detuning from resonance. For planar wavefronts, $k'_z(z) = 0$, and $\delta(t)$ corresponds to the familiar expression for the first-order Doppler shift (see Methods for details).

The experimental sequence is depicted in Fig. 1. We start in zone B by cooling all motional modes of the ion to $\bar{n} < 3$ using a combination of Doppler and electromagnetically induced transparency cooling[23], and then initialize the internal state by optical pumping into $|0\rangle$. The ion is then transported to zone A, and the laser beam used to implement the Hamiltonian is turned on in zone B. The ion is then transported through this laser beam to zone C. During the passage through the laser beam, we rapidly turn the beam off at time t_{off} and thus stop the qubit dynamics. The ion is then returned to the central zone B to perform state readout, which measures the qubit in the computational basis $\hat{\sigma}_z$ (for more details see Methods). The additional Hamiltonian $\hat{H}_s$ is implemented by offsetting the laser frequency used in the experiment by a detuning δ_L. For each setting of t_{off} and δ_L the experiment is repeated 100 times, allowing us to obtain an estimate for the qubit populations $\langle \hat{\sigma}_z^{meas}(t_{off}, \delta_L) \rangle$.

We first perform a comparison in which the ion velocity is the same but the beam position is changed. Thus we expect to obtain two different profiles for $\Omega(t)$ but the same velocity profile, which is closely related to $\delta(t)$. Experimental data is shown in Fig. 2 alongside the results of fitting performed using our iterative method. The beam positions used for each data set differ by $\sim 64\,\mu m$ along the transport axis, but the transport waveform used was identical. It can be seen from the residuals that the estimation is able to find a Hamiltonian, which results in a close match to the data.
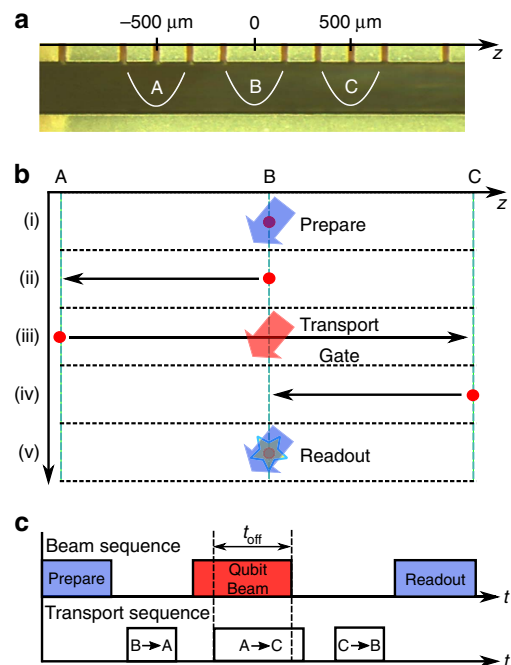
Figure 1 | Experimental sequence and timing. (**a**) The experiment is carried out in three zones of the trap indicated by A, B and C. (**b**) The experimental sequence involves steps (i) through (v). Preparation and readout are carried out on the static ion in zone B. The qubit evolves while the ion is transported through the laser beam in zone B in a transport operation taking the ion from zone A to zone C. (**c**) Experimental sequence showing the timing of applied laser beams and ion transport, including shutting off the laser beam during transport.
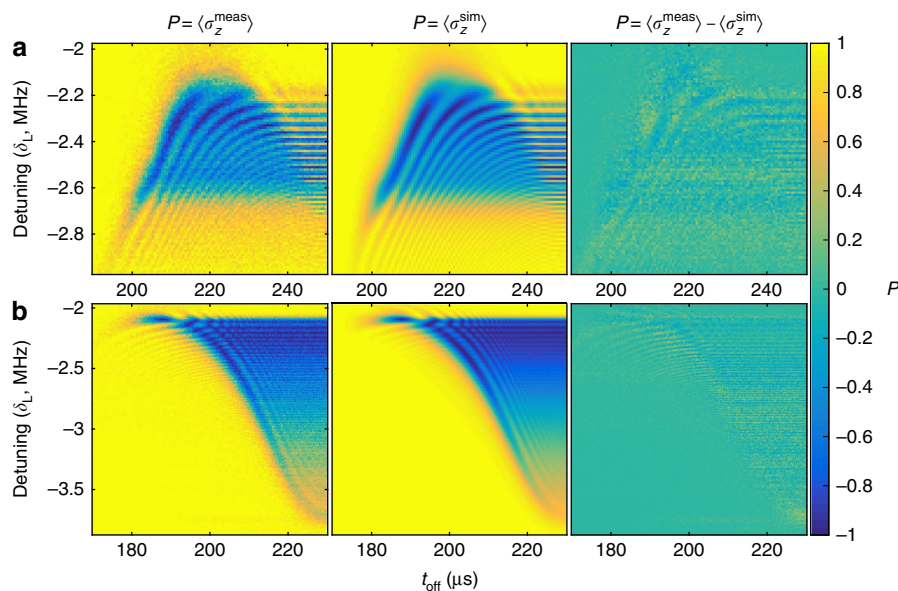
Figure 2 | Measured data, best fit and residuals. Spin population as a function of detuning and switch-off time of the laser beam. **a** is for a laser beam centred in zone B, while for **b** the beam was displaced towards zone C by $64\,\mu m$. From left to right are plots of the experimental data, the populations generated from the best fit Hamiltonian, and the residuals. Each data point results from 100 repetitions of the experimental sequence. The data in **a** consist of an array of 100×101 experimental settings, while that shown in **b** consists of an array of 201×201 settings. This leads to smaller error bars in the reconstructed Hamiltonian for the latter. For the Hamiltonian estimation, the data was weighted according to quantum projection noise.
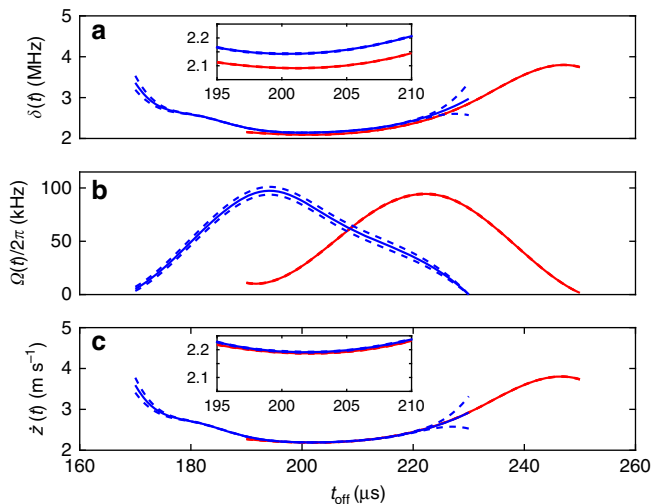
Figure 3 | Estimates of time-dependent co-efficients. (**a**) The effective detuning $\delta(t)$ and (**b**) Rabi frequency $\Omega(t)$ obtained from the two data sets. Blue and red solid lines show data obtained having the beam centred in zone B and with the beam displaced by a few tens of microns. Dashed lines indicate the s.e. on the mean of these estimates, which are obtained using resampling. For **a** the inset shows a close-up of the estimated $\delta(t)$ in the regions where the estimates overlap, showing that these do not give the same value. (**c**) The estimated velocity $\dot{z}(t)$ of the ion obtained after applying wavefront correction. The inset shows that this produces consistent results.

The estimated coefficients of the Hamiltonian extracted from the two data sets are shown in Fig. 3a,b. To estimate the relevant errors of our reconstruction, we have performed non-parametric resampling with replacement, optimizing for the solution using the same set of B-spline functions as was used for the experimental data to provide a new estimate for the Hamiltonian. This is repeated for a large number of samples, resulting in a

distribution for the estimated values of $\delta(t)$ and $\Omega(t)$ from which we extract statistical properties such as the s.e. The error bounds shown in Fig. 3 correspond to the s.e. on the mean obtained from these distributions (see Methods for further details). It can be seen that the values of $\delta(t)$ for the two different beam positions have a similar form but a fixed offset for the region where the reconstructions overlap. We believe that this effect arises from the non-planar wavefronts of the laser beam. Inverting the expression for $\delta(t)$ to obtain the velocity of the ion, we find $\dot{z}(t) = (\delta_L - \delta(t))/(k'_z(z)z + k_z(z))$. Using this correction, we find that the two velocity profiles agree if we assume that the ion passes through the centre of the beam at a distance of 2.27 mm before the minimum beam waist, a value which is consistent with experimental uncertainties due to beam propagation and possible mis-positioning of the ion trap with respect to the fixed final focusing lens. The velocity estimates taking account of this effect are shown in Fig. 3c.

Our second comparison involves using two different velocity profiles but with a common beam position. The resolution in both time and detuning were lower in this case than for the data shown in Fig. 2 (see Methods for the data). Figure 4 shows the results of the reconstruction. We observe that the estimated Rabi frequency profiles agree to within the error bars of the reconstruction. One interesting feature of this plot is that the error bars produced from the resampled data sets increase near the peak. We believe that this happens because the sampling time of the data is $0.5\,\mu s$, which starts to become comparable to the Rabi frequency (the Nyquist frequency is 1 MHz). To optimize the efficiency of our method, it would be advantageous to run the reconstruction method in parallel with data taking, thus allowing updating of the sampling time and frequency resolution based on the current estimates of parameter values.

Discussion
Our method for directly obtaining a non-commuting time-dependent Hamiltonian uses straightforward measurements of the qubit state in a fixed basis as a function of time and a
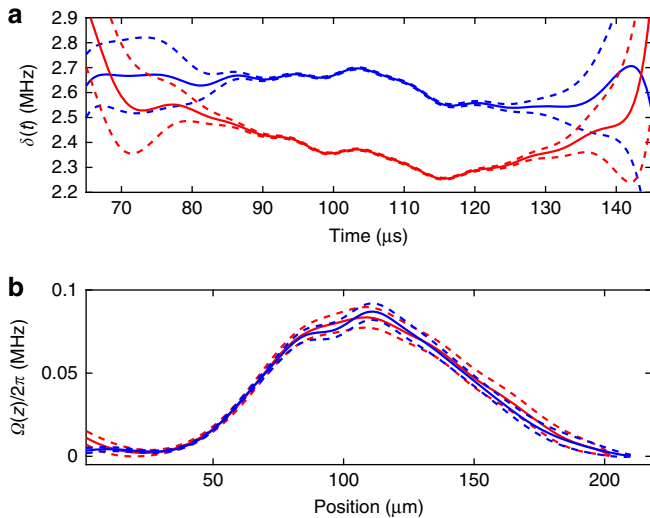
Figure 4 | Time-dependent detuning and spatial Rabi frequency.
(**a**) The estimated $\delta(t)$ obtained from the second pair of data sets (Fig. 8 in Methods). (**b**) The estimated Rabi frequency $\Omega(t)$ for the same two data sets. In each part, the blue and red solid lines show data obtained using different velocity profiles. Dashed lines indicate the s.e. on the mean of these estimates, which are obtained using resampling.

controlled offset to the Hamiltonian. Unlike schemes based on dynamical modulation or continuous strong driving, it avoids the need for control fields which act more strongly on the qubit than the Hamiltonian to be measured. This is a key advantage in quantum technologies where the Hamiltonian of interest is often already close to the limit of system drive strength. A process-tomography-based approach would require that for every time step multiple input states be introduced, and a measurement made in multiple bases[24–26]. This requires a much greater level of control than the method presented above. An effective modulation of the measurement basis arises in our approach due to the additional detuning δ_L. Nevertheless, it is also worth noting that tomography provides more information than our method: it makes no assumptions about the dynamics aside from that of a completely positive map while we require coherent dynamics. Extensions to our work are required to provide a rigorous estimation of the efficiency of the method in terms of the precision obtained for a given number of measurements, and to see whether a similar approach could be taken for non-unitary dynamics. We have recently used these methods to improve the control over the ion velocity, which is of direct value in optimizing transport operations in scalable trapped-ion quantum information processing[11,12,27], and will be essential for realizing multi-qubit transport gates[18]. We expect them to be applicable across a wide range of physical systems where such control is available, including those considered for quantum computation[4,6,7,28–31].

Methods

Derivation of Hamiltonian. The interaction of a laser beam with frequency ω_L and wave vector $\mathbf{k}(\mathbf{z}(t))$ with a two-level atom with resonant frequency ω_0 and time-dependent ion position $\mathbf{z}(t) = (0, 0, z(t))$ can be described in the Schrödinger picture by the Hamiltonian

$$\hat{H}_S = -\frac{\hbar\omega_0}{2}\hat{\sigma}_z - \hbar\Omega(z(t))\cos(\mathbf{k}(\mathbf{z}(t)) \cdot \mathbf{z}(t) - \omega_L t)\hat{\sigma}_x, \quad (5)$$

where the Rabi frequency $\Omega(z(t))$ gives the interaction strength between the laser and the two energy levels. We can define the laser phase at the position of the ion as $\Phi(t) = \phi(t) - \omega_L t$ with $\phi(t) = \mathbf{k}(\mathbf{z}(t)) \cdot \mathbf{z}(t) = k_z(z(t))z(t)$ and $k_z(z(t)) = |\mathbf{k}|\cos(\theta(t))$ being the projection of the laser beam onto the z-axis along which the ion is transported. Here $\theta(t)$ is the angle between the wave vector $\mathbf{k}(\mathbf{z}(t))$ and the transport axis evaluated at position $z(t)$. Moving to a rotating frame using the
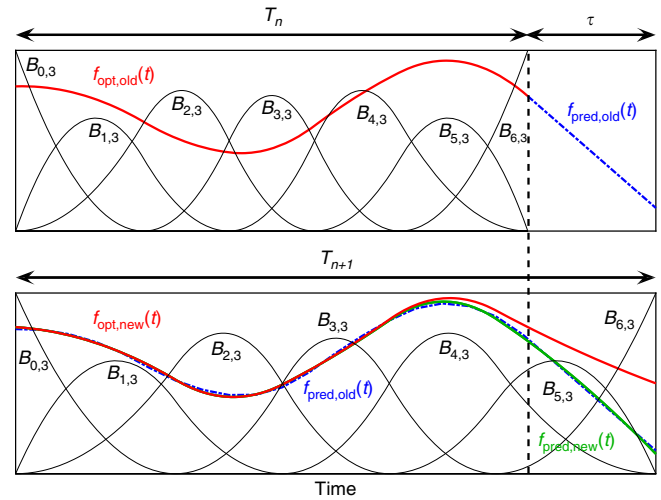
Figure 5 | Extending the horizon estimation. The steps performed when extending the time horizon from T_n to T_{n+1} are illustrated. We first predict in the old basis, then move to the new basis, and finally optimize again. The figure also shows the B-splines $B_{i,k}(t)$.

unitary transformation $U = e^{-i\frac{\Phi(t)}{2}}$ and applying the rotating wave approximation with respect to optical frequencies, we obtain

$$\hat{H}_I = \frac{\hbar}{2}\left(-\Omega(t)\hat{\sigma}_x + \left(-\omega_0 - \dot{\Phi}(t)\right)\hat{\sigma}_z\right). \quad (6)$$

Defining a static detuning $\delta_L = \omega_L - \omega_0$, we obtain

$$\hat{H}_I = \frac{\hbar}{2}\left(-\Omega(t)\hat{\sigma}_x + \left(\delta_L - \dot{\phi}(t)\right)\hat{\sigma}_z\right). \quad (7)$$

with

$$\delta(t) = \delta_L - \dot{\phi}(t), \quad (8)$$

which is the expression used in the main text.

B-spline curves and optimization algorithm. The set of polynomial B-spline functions $B_{i,k}(t)$ of order k are recursively defined over the index i over a set of points $\mathbf{K} = \{t_0, t_1, .., t_{n+k}\}$, which is referred to as the knot vector[20].

$$
\begin{aligned}
B_{i,1}(t) &= \begin{cases} 1 & t_i \le t \le t_{i+1} \\ 0 & \text{otherwise} \end{cases} \\
B_{i,k}(t) &= \omega_{i,k}(t)B_{i,k-1}(t) + \left(1 - \omega_{i+1,k}(t)\right)B_{i+1,k-1}(t). \\
\omega_{i,k}(t) &= \begin{cases} \frac{t - t_i}{t_{i+k-1} - t_i} & \text{if } t_i \ne t_{i+k-1} \\ 0 & \text{otherwise} \end{cases}
\end{aligned}
\quad (9)
$$

Figure 5 gives a visualization of the B-splines $B_{i,k}(t)$ and a B-spline curve. The B-spline construction ensures that any linear combination of the B-splines is continuous and has $(k - 2)$ continuous derivatives. The knot vector $\mathbf{K}$ determines how the basis functions are positioned within the interval $[t_0, t_{n+k}]$. We notice that for our Hamiltonian the spacing of the B-splines is not critical, which we think is due to the smoothness of the variations in our Hamiltonian parameters $\delta(t)$ and $\Omega(t)$. We therefore used the Matlab function spap2 to automatically choose a suitable knot vector and restricted ourselves to optimizing the coefficients α_i. We collect all coefficients α_i for $\delta(t)$ and $\Omega(t)$ and store them in a single vector $\boldsymbol{\alpha}$.

A detailed algorithmic summary of our implementation of the EHE method is given below.

1. Searching for a starting point: here we reconstruct the Hamiltonian for a first, minimal time horizon such that we can then use this as a starting point to iteratively extend the horizon as described in step 2.

   (a) Choose an initial time horizon such that it contains the region where the first discernible qubit dynamics occur.
   (b) Cut down the number of fitting parameters as much as possible, for example, by using few B-splines of low order. This amounts to choosing empirically a low number of B-splines (and thus the length of $\boldsymbol{\alpha}_0$), which might represent $\delta(t)$ and $\Omega(t)$ over the given region.
   (c) Use a non-linear least-squares fitting routine to minimize J by varying the parameters $\boldsymbol{\alpha}_0$. In the case that the initial fit is poor or no minimum is found, try new initial conditions, change the number of B-spline functions, or manually adjust the function using prior knowledge of the physical system under consideration.
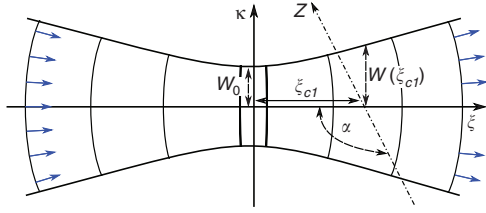
Figure 6 | Beam and ion transport. The beam propagation direction lies along the ξ-axis and the ion is transported along the z-axis lying on the $\kappa\xi$-plane as indicated. Normalized vectors representing $\mathbf{e}_l(\kappa, \xi)$ lying perpendicular to the wavefronts are indicated by the blue arrows.

This procedure is used to provide a starting point for the optimization over the initially chosen window, which is typically performed with a set of higher order B-splines. From this starting point, we iteratively extend the fitting method to the full data set as follows.

2. Extend the horizon: this step is repeated until the whole time horizon is covered. It consists of the following sequence, which is illustrated in Fig. 5.

   (a) Extend the time horizon by τ from T_n to $T_{n+1} = T_n + \tau$.
   (b) Extrapolate $f_{\mathrm{opt,old}}(t)$ within τ, for example, using fnxtr in Matlab.
   (c) Adapt the B-splines to the new time horizon T_{n+1} and represent $f_{\mathrm{pred,old}}(t)$ in the new basis, giving $f_{\mathrm{pred,new}}(t)$. In Matlab one can use spap2 to do this.
   (d) Use $f_{\mathrm{pred,new}}(t)$ as the initial guess for a weighted non-linear least-squares fit over the extended time span up to T_{n+1}.
   (e) Judge the results of the fit based on its reduced χ^2-value χ^2_{red}. If it is below a specified bound, continue with an additional iteration of steps a–d, repeating until the full region of the data is covered. Otherwise, attempt the following fall-back procedures:

      i Reduce τ, the time by which the time horizon is extended, and try again.
      ii Increase the number of B-splines and try again.
      iii Try again using a different starting point.

If all these fail, we have to resort to increasing the bound on χ^2_{red}.

3. Post-processing: the following steps are optional and were performed manually in cases where we wished to improve the fit or examine its behaviour.

   (a) The optimization over the whole time horizon was re-run using different numbers of B-splines for $\delta(t)$ and $\Omega(t)$. This was used to check the sensitivity of the fit.
   (b) The optimization over the whole time horizon was re-run using a starting point based on the previously found optimum plus randomized deviations. This tested the robustness of the final fit.

Wavefront correction. For plane waves, we find that $\dot{\phi}(t) = \mathbf{k} \cdot \mathbf{v}(t)$, which is the well-known expression for the first-order Doppler shift. For transport through a real Gaussian beam, the wave vector direction changes with position. Taking this into account, the derivative of $\phi(t)$ becomes

$$\dot{\phi}(t) = \left[k_z'(z(t))z(t) + k_z(z(t)) \right] \dot{z}(t) \qquad (10)$$

where $k_z' = \mathrm{d}k_z/\mathrm{d}z$ and $\dot{z}(t)$ is the component of the ion velocity along the z-axis. We extract $\delta(t)$ using our Hamiltonian estimation procedure, thus to obtain the velocity of the ion we use

$$\dot{z}(t) = \frac{-\delta(t) + \delta_{\mathrm{L}}}{k_z'(z(t))z(t) + k_z(z(t))}. \qquad (11)$$

As the ion moves through the beam it experiences the same magnitude of the wave vector $|\mathbf{k}| = 2\pi/\lambda$, but the angle θ between the ion direction and the wave vector changes. Written as a function of this angle, the velocity becomes

$$\dot{z}(t) = \frac{-\delta(t) + \delta_{\mathrm{L}}}{-|\mathbf{k}|\sin(\theta(z(t)))\theta'(z(t))z(t) + |\mathbf{k}|\cos(\theta(z(t)))} \qquad (12)$$

where $\theta'(z(t)) = \mathrm{d}\theta(z(t))/\mathrm{d}z(t)$. We parameterize our Gaussian beam according to Fig. 6. The phase is given as a function of both the position along the beam axis ξ
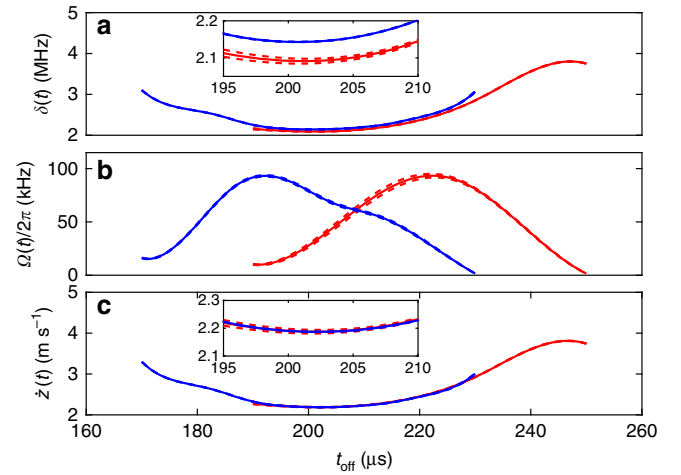
Figure 7 | Parametric bootstrap resampling. Predictions for the effective detuning $\delta(t)$ in **a**, Rabi frequency $\Omega(t)$ in **b** and velocity $\dot{z}(t)$ in **c**. Blue and red solid lines show data obtained having the beam centred in zone B and with the beam displaced by a few tens of micron. Dashed lines represent the s.e. on the mean of these estimates obtained using parametric bootstrap resampling, assuming quantum projection noise. This can be compared with the error bounds obtained from the non-parametric method, which are shown in Fig. 3 in the main text. The bounds are tighter for the parametric bootstrapping.

and the perpendicular distance from this axis κ by[32]

$$\varphi(\kappa, \xi) = |\mathbf{k}|\xi - \zeta(\xi) + \frac{|\mathbf{k}|\kappa^2}{2R(\xi)}. \qquad (13)$$

where the Gaussian beam parameters include the beam waist $W(\xi)$, the radius of curvature $R(\xi)$, the Rayleigh range ξ_{R} and the Guoy phase shift $\zeta(\xi)$. These are given by the expressions

$$
\begin{aligned}
W(\xi) &= W_0\sqrt{1 + \left(\frac{\xi}{\xi_{\mathrm{R}}}\right)^2} \\
R(\xi) &= \xi\left(1 + \left(\frac{\xi_{\mathrm{R}}}{\xi}\right)^2\right) \\
\zeta(\xi) &= \tan^{-1}\left(\frac{\xi}{\xi_{\mathrm{R}}}\right) \\
\xi_{\mathrm{R}} &= \frac{\pi W_0^2}{\lambda} \\
k &= \frac{2\pi}{\lambda}
\end{aligned}
\qquad (14)
$$

where W_0 is the minimum beam waist and λ the laser wavelength. The ion moves along the z-axis as shown in Fig. 6. In the $\kappa\xi$-plane a unit vector $\mathbf{e}_l(\kappa, \xi)$ perpendicular to the wavefronts is given by

$$\mathbf{e}_l(\kappa, \xi) = \frac{\nabla\varphi(\kappa, \xi)}{\|\nabla\varphi(\kappa, \xi)\|} \qquad (15)$$

and the unit vector $\mathbf{e}_v$ pointing along the direction of transport is given by

$$\mathbf{e}_v = \begin{bmatrix} \cos(\alpha) \\ \sin(\alpha) \end{bmatrix} \qquad (16)$$

The angle $\theta(\xi)$ between the wave and position vector is then given by the dot product

$$\theta(\kappa) = \cos^{-1}(\mathbf{e}_n \cdot \mathbf{e}_v). \qquad (17)$$

which can be written in terms of the full set of parameters above as

$$
\begin{aligned}
\theta(\kappa) &= \cos^{-1}(\gamma_1 + \gamma_2) \\
\gamma_1 &= \frac{\cos(\alpha)\left(-2\xi_{\mathrm{R}}\left(\xi^2 + \xi_{\mathrm{R}}^2\right) + k\kappa^2\left(\xi_{\mathrm{R}}^2 - \xi^2\right) + 2k\left(\xi^2 + \xi_{\mathrm{R}}^2\right)^2\right)}{\eta(\kappa)} \\
\gamma_2 &= \frac{\sin(\alpha)2k\kappa\xi\left(\xi^2 + \xi_{\mathrm{R}}^2\right)}{\eta(\kappa)} \\
\eta(\kappa, \xi) &= \left(\xi^2 + \xi_{\mathrm{R}}^2\right)\left[4\left(\frac{k\kappa\xi}{\xi^2 + \xi_{\mathrm{R}}^2}\right)^2 + \left(-\frac{2\xi_{\mathrm{R}}}{\xi^2 + \xi_{\mathrm{R}}^2} + k\left(2 + \frac{\kappa^2\left(\xi_{\mathrm{R}}^2 - \xi^2\right)}{\left(\xi^2 + \xi_{\mathrm{R}}^2\right)^2}\right)\right)^2\right] \\
\kappa(t) &= z(t)\sin(\alpha)
\end{aligned}
\qquad (18)
$$

where in our experiments $\alpha = 3\pi/4$.

Using equations (12) and (18), we examined the value of ξ_{cl} required for the velocity to match for our two beam positions. We find that they agree for $\xi_{\mathrm{cl}} = -2.27$ mm, which is within the experimental uncertainties for our set-up.
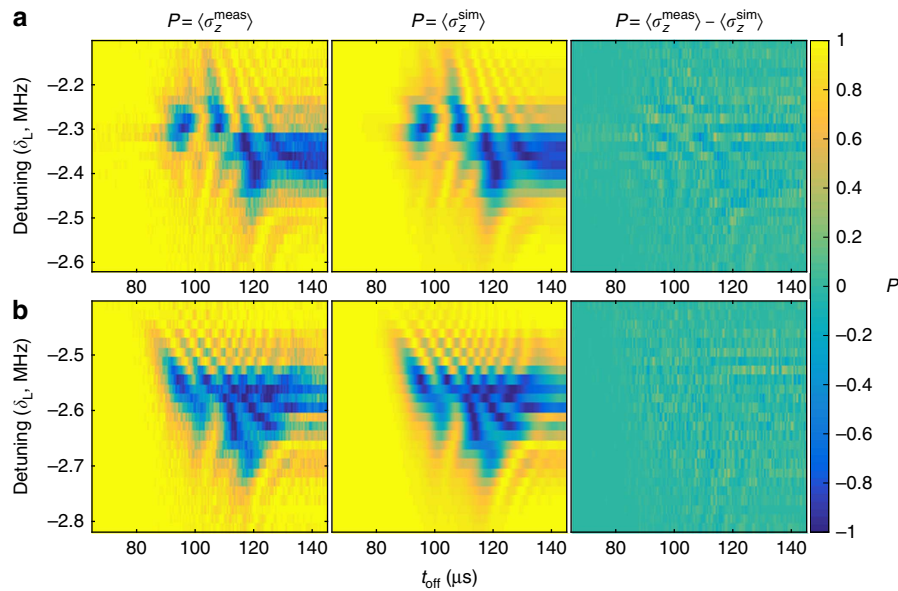
Figure 8 | Measured data, estimation and residuals. Spin population as a function of detuning and switch-off time of the laser beam, for the data sets used to obtain the reconstructed parameters shown in Fig. 4. **a** corresponds to the blue reconstruction while **b** to the red one. Each data point results from 100 repetitions of the experimental sequence. For the Hamiltonian estimation, the data was weighted according to quantum projection noise.

Error estimation. To estimate the errors of the time-dependent functions, we use non-parametric bootstrapping[33]. The process is summarized as follows:

1. Estimate initial solution: estimate the time-dependent functions from the original data using Hamiltonian estimation.
2. Resampling: create N_s sample solutions for all time-dependent functions in the following way:

   (a) Form a sample set by randomly picking with replacement from the photon count data used in qubit detection.
   (b) Re-estimate new time-dependent functions by optimizing over the full time span, using the solution found in (1) as a starting point.
   (c) Record the reduced χ^2-values $\chi^2_{red,r}$ for each sample r along with the B-spline curve coefficients $\boldsymbol{\alpha}_r$.

3. Post-process samples:

   (a) Form a histogram of the χ^2-values $\chi^2_{red,r}$.
   (b) Find and fit a normal-like distribution to the histogram with preference to the spread with lowest lying $\chi^2_{red,r}$ in the case of a multi-modal distribution. From the fit obtain the mean reduced-χ^2-value $\langle\chi^2_{red,r}\rangle$, as well as the s.d. σ_χ.
   (c) Eliminate the outlier samples by removing all $\boldsymbol{\alpha}_r$ with $\chi^2_{red,r}$ values that are 3–5σ_χ from the mean $\langle\chi^2_{red,r}\rangle$.
   (d) Form a matrix $\mathbf{Y}$, where each row vector is a sample set of coefficients $\boldsymbol{\alpha}_r$ that remained after step 3(c).

4. Obtain statistics:

   (a) Find the mean B-spline coefficients $\langle\alpha\rangle$ of equation (4) by taking the mean over the column vectors of $\mathbf{Y}$ with each element of the mean given by $\langle\alpha\rangle_i=\langle\alpha_i\rangle$.
   (b) Find the covariance matrix $\Sigma=\text{cov}(\mathbf{Y}^\alpha)$ with $\Sigma_{ij}=\text{E}[(\alpha_i-\langle\alpha_i\rangle)(\alpha_j-\langle\alpha_j\rangle)]$ with E the expectation operator. The s.d. of each of the mean coefficients $\langle\alpha_i\rangle$ is given by $\sigma_{\langle\alpha_i\rangle}=\sqrt{\Sigma_{ii}}$. We record these values in a row vector $\boldsymbol{\sigma}_{\langle\alpha_i\rangle}$.

In evaluating errors using bootstrapping, we use the same set of spline polynomials as were used for the final optimization stage in the data, which makes the reconstruction more reliable in converging to a minimum. We thus expect that the parameter space explored in evaluating the errors is not the same as for the *ab initio* estimation of the Hamiltonian. This is apparent in the regions of the data where the final estimate has large error bounds (for example, in Fig. 4), where an even larger spread might be expected (the dynamics has stopped evolving at this point). We think that the net effect is to under-estimate the errors in the regions where the Hamiltonian is uncertain, but that the error bars given in the central region (where the Hamiltonian is well-defined) are close to what would be obtained through a full optimization.

We have also applied parametric bootstrapping to obtain the error bounds shown in Fig. 7. The difference to the non-parametric case is that in point (2) the samples are created using the solutions obtained from (1) and adding quantum projection noise. For each sample the Hamiltonian is estimated. The estimates from multiple samples are used to construct error bounds in the same manner as for the non-parametric resampling. We have found that the error bounds obtained from parametric bootstrapping are lower compared with that of the non-parametric case as shown in Fig. 3. We think this is due to the latter exploring deviations around a single minimum in the optimization landscape, while the case resampling arrives at different local minima, which are spread over a wider region.

Single beam profile with two different velocity profiles. To verify that our method can also consistently estimate the Rabi frequency profile, we measure a second pair of data sets in which we take two different velocity profiles using the same beam position. This data is shown in Fig. 8. Also shown are the best-fits obtained from the reconstructed Hamiltonians. The parameter variations obtained from the reconstructed Hamiltonians for these data sets can be found in the main text in Fig. 4. The sampling rate of the data in these data sets was 2 MHz, resulting in a Nyquist frequency of 1 MHz.

References

1. Cole, J. H. et al. Identifying an experimental two-state hamiltonian to arbitrary accuracy. *Phys. Rev. A* **71,** 062312 (2005).
2. Mitra, A. & Rabitz, H. Identifying mechanisms in the control of quantum dynamics through hamiltonian encoding. *Phys. Rev. A* **67,** 033407 (2003).
3. de Castro, R. R. & Rabitz, H. Laboratory implementation of quantum-control-mechanism identification through hamiltonian encoding and observable decoding. *Phys. Rev. A* **81,** 063422 (2010).
4. Zhang, J. & Sarovar, M. Quantum hamiltonian identification from measurement time traces. *Phys. Rev. Lett.* **113,** 080401 (2014).
5. de Castro, R. R., Cabrera, R., Bondar, D. I. & Rabitz, H. Time-resolved quantum process tomography using hamiltonian-encoding and observable-decoding. *N. J. Phys.* **15,** 025032 (2013).
6. Devitt, S. J., Cole, J. H. & Hollenberg, L. C. L. Scheme for direct measurement of a general two-qubit hamiltonian. *Phys. Rev. A* **73,** 052317 (2006).
7. Shabani, A., Mohseni, M., Lloyd, S., Kosut, R. L. & Rabitz, H. Estimation of many-body quantum hamiltonians via compressive sensing. *Phys. Rev. A* **84,** 012107 (2011).
8. Zhao, N. et al. Sensing single remote nuclear spins. *Nature Nanotech.* **7,** 657–662 (2012).

9. Blais, A., Huang, R.-S., Wallraff, A., Girvin, S. M. & Schoelkopf, R. J. Cavity quantum electrodynamics for superconducting electrical circuits: an architecture for quantum computation. *Phys. Rev. A* **69,** 062320 (2004).
10. Devoret, M. H. & Schoelkopf, R. J. Superconducting circuits for quantum information: an outlook. *Science* **339,** 1169–1174 (2013).
11. Bowler, R. *et al.* Coherent diabatic ion transport and separation in a multizone trap array. *Phys. Rev. Lett.* **109,** 080502 (2012).
12. Walther, A. *et al.* Controlling fast transport of cold trapped ions. *Phys. Rev. Lett.* **109,** 080501 (2012).
13. Kotler, S., Akerman, N., Glickman, Y. & Ozeri, R. Non-linear single-spin spectrum analyzer. *Phys. Rev. Lett.* **110,** 110503 (2013).
14. Magesan, E., Cooper, A., Yum, H. & Cappellaro, P. Reconstructing the profile of time-varying magnetic fields with quantum sensors. *Phys. Rev. A* **88,** 032107 (2013).
15. Cooper, A., Magesan, E., Yum, H. & Cappellaro, P. Time-resolved magnetic sensing with electronic spins in diamond. *Nat. Commun.* **5,** 3141 (2014).
16. Barnes, E. & Das Sarma, S. Analytically solvable driven time-dependent two-level quantum systems. *Phys. Rev. Lett.* **109,** 060401 (2012).
17. Barnes, E. Analytically solvable two-level quantum systems and landau-zener interferometry. *Phys. Rev. A* **88,** 013818 (2013).
18. Leibfried, D., Knill, E., Ospelkaus, C. & Wineland, D. J. Transport quantum logic gates for trapped ions. *Phys. Rev. A* **76,** 032324 (2007).
19. de Clercq, L. *et al.* Parallel Transport-Based Quantum Logic Gates With Trapped Ions. *Phys. Rev. Lett.* **116,** 080502 (2016).
20. de Boor, C. *A Practical Guide to Splines (Applied Mathematical Sciences),* revised edn (Springer, 2001).
21. Muske, K. R. & Rawlings, J. B. in *Methods of Model Based Process Control* (ed. Berber, R.) (Springer, 1995).
22. Rao, C. V., Rawlings, J. B. & Mayne, D. Q. Constrained state estimation for non-linear discrete-time systems: stability and moving horizon approximations. *IEEE Trans. Automat. Contr.* **48,** 246–258 (2003).
23. Roos, C. F. *et al.* Experimental demonstration of ground state laser cooling with electromagnetically induced transparency. *Phys. Rev. Lett.* **85,** 5547 (2000).
24. Chuang, I. L. & Nielsen, M. A. Prescription for experimental determination of the dynamics of a quantum black box. *J. Mod. Opt.* **44,** 2455–2467 (1997).
25. Poyatos, J. F., Cirac, J. I. & Zoller, P. Complete characterization of a quantum process: the two-bit quantum gate. *Phys. Rev. Lett.* **78,** 390 (1997).
26. Riebe, M. *et al.* Process tomography of ion trap quantum gates. *Phys. Rev. Lett.* **97,** 220407 (2006).
27. Wineland, D. J. *et al.* Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *J. Res. Natl. Inst. Stand. Technol.* **103,** 259–328 (1998).
28. Martinis, J. M. & Geller, M. R. Fast adiabatic gates using only σ_z control. *Phys. Rev. A* **90,** 022307 (2014).
29. Schutjens, R., Abu Dagga, F., Egger, D. J. & Wilhelm, F. K. Single-qubit gates in frequency-crowded transmon systems. *Phys. Rev. A* **88,** 052330 (2013).
30. Barends, R. *et al.* Digital quantum simulation of fermionic models with a superconducting circuit. *Nat. Commun.* **6,** 7654 (2015).
31. Bonato, C. *et al.* Optimized quantum sensing with a single electron spin using real-time adaptive measurements. *Nature Nanotech* **11,** 247–252 (2016).
32. Saleh, B. E. A. & Teich., M. C. *Fundamentals of Photonics* (Wiley, 2007).
33. Murphy., K. P. *Machine Learning: a Probabilistic Perspective* (MIT press, 2012).

Acknowledgements

Author contributions

Experimental data was taken by L.E.d.C., R.O., M.M. and D.N. using apparatus built up by all authors. Data analysis was performed by L.E.d.C. and R.O. The paper was written by J.P.H., L.E.d.C. and R.O., with input from all authors. The work was conceived by J.P.H. and L.E.d.C.

Additional information