

Machine-learning optimal control pulses in an optical quantum memory experiment

Elizabeth Robertson^{1,2,*}, Luisa Esguerra^{1,2}, Leon Meßner¹, Guillermo Gallego^{3,4,5,6} and Janik Wolters^{1,2,3}

¹*Deutsches Zentrum für Luft- und Raumfahrt e.V. (DLR), Rutherfordstraße 2, 12489 Berlin, Germany*


²*Technische Universität Berlin, Institute for Optics and Atomic Physics, Hardenbergstraße 36, 10623 Berlin, Germany*

³*Technical University of Berlin, Straße des 17. Juni 135, 10623 Germany*

⁴*Einstein Center Digital Future, Wilhelmstraße 67, Berlin 10117*

⁵*Science of Intelligence Excellence Cluster, Marchstraße 23, Berlin 10587*

⁶*Robotics Institute Germany, Straße des 17. Juni 135, Berlin 10623*

 (Received 10 January 2024; revised 27 May 2024; accepted 8 July 2024; published 8 August 2024)

Efficient optical quantum memories are a milestone required for several quantum technologies, including repeater-based quantum key distribution and on-demand multiphoton generation. We present an efficiency optimization of an optical electromagnetically induced transparency (EIT) memory experiment in a warm cesium vapor using a genetic algorithm and analyze the resulting wave forms. The control pulse is represented either as a Gaussian or free-form pulse and the results from the optimization are compared. We see an improvement factor of 3(7)% when using optimized free-form pulses. By limiting the allowed pulse energy in a solution, we show an energy-based optimization giving a 30% reduction in energy, with minimal efficiency loss.

DOI: [10.1103/PhysRevApplied.22.024026](https://doi.org/10.1103/PhysRevApplied.22.024026)

I. INTRODUCTION

Optical memories have long been recognized as a required technology in the implementation of different quantum protocols, notably on-demand multiphoton generation [1], repeater-based quantum key distribution (QKD) [2–5], and the translation of flying to stationary qubits, among others [6,7]. Indeed, the aforementioned protocols have highly efficient operation as a requirement and often the memory efficiency bounds the attainable end-to-end efficiency of the system. In, e.g., QKD implementations, this leads to a logarithmic scaling in the time required for entanglement distribution [3]. Moreover, optical memories find applications in classical analogue computing, in systems such as reservoir computing, where the efficiency of the optical memory provides a bound on the memory capacity of the reservoir [8].

Similarly to the variety one sees in the “qubits zoo,” quantum memories are equally diverse in their form [9]. Spanning from single atoms [10] to solid-state devices [11,12], a range of optical quantum memories [13] have been investigated, providing a wide range of efficiencies. The most efficient demonstrated memory systems are those operating at the ultracold regimes. By removing

several sources of decoherence, efficiencies of 68%–92% [14–20] have been achieved. However, solid-state devices also boast high-efficiency operation, shown between 56% and 69% [21,22]. Warm-vapor atomic memories based on electromagnetically induced transparency (EIT) have been highlighted as one of the most promising of these systems, as they are technologically simple, can be multiplexed [23], and have been modeled to have high operating efficiency [24]. Using strong magnetic control of atomic ensembles, an 82% internal efficiency at room temperature has been demonstrated [25]. Moreover, warm-vapor memories have been shown to have an acceptance bandwidth of 0.66 GHz, making them suitable for interfacing with semiconductor quantum dots and other single-photon sources [26].

A theoretical formulation and optimization of the three-level Λ EIT scheme in free space has been presented in Refs. [24,27]. There, the authors used a gradient-ascent approach to optimize the control field, using an analytically calculable gradient derived from the atomic spin wave, to learn the temporal shape of the optimal optical pulse. These simulation-optimized pulses were then transferred to the experiment, where they were shown to perform well [28,29]. However, the model presented in these works is limited in the physics it accounts for; indeed, the presence of a fourth level changes the shape of the optimal control

*Contact author: elizabeth.robertson@dlr.de

pulses, as shown in Ref. [30]. Moreover, there are further experimental effects such as four-wave mixing (FWM) and the inhomogeneous broadening of the excited state that are not accounted for in the four-level system models and the optimal operating conditions of a memory experiment.

Using the genetic algorithm, we learn the optimal control pulses of a warm-cesium-vapor quantum memory, where the temporal shape is encoded either as a Gaussian, with the amplitude, pulse width, and delay as free parameters, or as a 16-parameter free-form pulse. In this paper, we apply an optimization process to the experiment as a whole. In-experiment optimizations boast the benefit of accounting for a variety of experimental effects; device-specific transfer functions of optical modulators, intersystem delays between signal and control, and physical effects not captured in model systems. However, as the spin-wave gradients are not accessible in an experimental setting, we must consider an alternative gradient-free learning approach. Genetic algorithms are widely acknowledged as noise-robust gradient-free optimization algorithms, making them attractive for use in a wide range of atomic (optics) experiments [31–34].

We find that the efficiency of memory experiments carried out with Gaussian pulses is similar to those of free-form pulses, in the device-resolvable range of the experiment. This is confirmed in theory, albeit for the storage of shorter signals [35]. We show experimentally that this trend holds for larger pulse widths; i.e., we find that the use of free-form pulses provides an average improvement of 3(7)% for signals ranging between $\tau_{\text{FWHM}}\gamma = \{1.9, 21.5\}$. Here, τ_{FWHM} is the full width at half maximum (FWHM) of the signal pulse, given in units of τ , where τ is $1/\gamma$ and γ is the excited state decay rate. Moreover, we demonstrate that the temporal regions of particular importance to the efficiency when using free-form pulses are similar to those temporally overlapping with the signal pulse [36]. We note while this approach

is platform agnostic, our choice of temporal signal length means that this work only considers the operating regime of the EIT protocol. We also illustrate the possibility of learning optimal pulses under further objective constraints, such as total-pulse-energy minimization. Here, we show that we can reduce the energy of the learned signal and control pulses by 30% with a minimal trade-off of 4(6)% in efficiency. This finding might have important implications for reducing the readout noise—which is a well-known issue in warm-vapor memories.

This work is structured as follows. In Sec. II, we first present the optical memory (Sec. II A) and then the genetic algorithm (Sec. II B). The analysis of the results is divided into two sections. First, we consider the optimization of stored signals of different width (Sec. III). Then, in Sec. III B, we discuss the results of the energy optimization. A brief discussion of improvements to the method and an outlook for further research are given in Sec. IV.

II. OPTICAL QUANTUM MEMORY OPTIMIZATION

A. Warm-vapor memory

The experimental setup that we optimize has been described extensively in Ref. [37]; in Fig. 1(a), we depict a compressed setup and below we emphasize the differences in the setup used in this work.

We use two lasers with linear orthogonal polarization, which are offset locked to a frequency difference of 9.2 GHz, as the signal (S) and control (C) lasers. The signal and control pulses are generated with an electro-optic modulator and an acousto-optic modulator, respectively. This enables us to reach peak pulse powers of $P_{\text{C,cw}} = 12.9$ mW, without the semiconductor optical amplifier and spontaneous emission filtering system used in Ref. [37]. A pump laser with power $P_{\text{P,cw}} = 15.1$ mW is locked on the $F = 4 \rightarrow F' = 4$ transition and counter-propagates to




FIG. 1. The experimental schematics. (a) An overview of the online genetic algorithm optimization setup. Possible solutions (control pulses) are generated by the genetic algorithm and are evaluated by carrying out a memory experiment. EOM, electro-optic modulator; AOM, acousto-optic modulator; AFG, arbitrary function generator; CP1, CP2, calcite prisms; PD, photodiode; Cs, cesium-vapor cell; magn. shield, magnetic shield. (b) The three-level Λ system used in the memory protocol. The signal (control) lasers are red detuned by 1 GHz from their resonant hyperfine transitions. A resonant pump laser populates the ground state $|g\rangle$. (c) A typical memory experiment sequence. We denote $t_0 = 0$ as the center of the signal pulse to be stored and we read out such that the time between the signal in and the control read pulse is $\Delta t = 200$ ns.

the signal and control lasers. We red detune both lasers to be 1 GHz off the atomic transition; a level scheme indicating the laser frequencies is shown in Fig. 1(b). The wavelength of the control laser is monitored using a wavemeter and is locked to it using a simple feedback loop, to prevent frequency drifts during the algorithm evolution. For each memory experiment carried out during the genetic algorithm evolution, we pump for 400 ns and when evaluating the final efficiency, we pump for 10 μ s. Differing to Ref. [37], we operate the coherent signal pulses well above single-photon level, at a continuous wave power of 126 μ W. Moreover, we detect the output of the memory directly on an amplified photodiode; i.e., without further filtering.

A memory experiment is illustrated in Fig. 1(c). The write pulse (purple) mediates the storage of the input signal pulse (blue) in a collective spin state of cesium atoms. After a time $\Delta t = 200$ ns, the read pulse (red) retrieves the stored optical field from the collective spin state; this output signal is referred to as the retrieved pulse. Any component of the optical field that is not stored in this process is referred to as the leakage. In this work, we focus on learning the shape of the write pulse: the read pulse is fixed as a Gaussian with an FWHM of 40 ns. To quantify the performance of the memory, we define an internal memory efficiency η_{int} :

$$\eta_{\text{int}} = \frac{\int_{\Delta t - \Delta t/2}^{\Delta t + \Delta t/2} |E_{\text{out}}|^2 dt}{\int_{-\Delta t/2}^{\Delta t/2} |E_{\text{in}}|^2 dt}, \quad (1)$$

where E_{out} is the retrieved signal pulse, E_{in} is the initial signal, and Δt is the time between the initial signal and the retrieval control pulse. This input pulse, $|E_{\text{in}}|^2$, is measured using the same detector, PD1 (Thorlabs APD0815), in the far-detuned regime, such that the optical losses in the system remain the same for both measurements but the atoms do not absorb the signal. We red detune the lasers by $\Delta = 2$ GHz for this measurement. All the efficiencies of the final learned pulse are evaluated in this way. The efficiencies in this work are reported for a storage time of $\Delta t = 200$ ns and are not extrapolated back to $\Delta t = 0$ as would be standard in memory-systems comparison. Consequently, we expect that we are under-representing the efficiencies by a factor of about 1.3.

To ensure consistency across multiple experimental runs, we implement corrections for various sources of drift. The primary sources of drift or noise are identified as originating from the control laser frequency, the frequency offset between the signal and control lasers, the bias of the electro-optic modulator (EOM), and the temperature stabilization of the atoms. Certain experimental parameters, such as the frequency offset between the control and signal lasers, are stabilized using continuous feedback signals, while most noise sources in the system are managed

using an epoch-wise feedback loop. In this setup, the system is reset to target values at the beginning of each epoch and is assumed to remain stable throughout its duration until after measurement of the memory experiments, where it is then checked and readjusted. At the start of each epoch, we measure the control laser frequency, correcting it if it deviates more than 60 MHz from the set point. Subsequently, we set the bias voltage of the EOM to minimum transmission. During the evaluation of the fitness function (execution of the memory experiment), heating of the atomic cell is deactivated to minimize stray magnetic fields. Following measurement, the cell is reheated to 60 $^{\circ}$ C before the execution of the algorithm is resumed.

B. Genetic algorithm

The genetic algorithm [38] was developed by considering how adaptation can be imported to computer systems and thus to objective optimization [39]. The goal of a genetic algorithm is to iteratively find solutions that fulfill the objective of a function well.

In Fig. 2, we show a flow chart of a genetic algorithm. Briefly, solutions (individuals) are encoded using parameters (genes). Each solution is evaluated using an objective (fitness) function, which gives the optimization goal. Solutions that have a high fitness, i.e., are appropriate solutions to the objective, are selected to form the next generation. The genes of the selected solution are crossed over and mutated, generating the population of the next generation. The genetic algorithm is implemented using the PYTHON package PyGAD [40].

Here, we wish to find the optimal control-pulse shape, which maximizes the internal memory efficiency. Previously, Gaussians have been demonstrated as good approximations of an ideal control pulse and they provide a good benchmark for performance analysis against a more general free-form pulse [36,37]. To be able to distinguish the

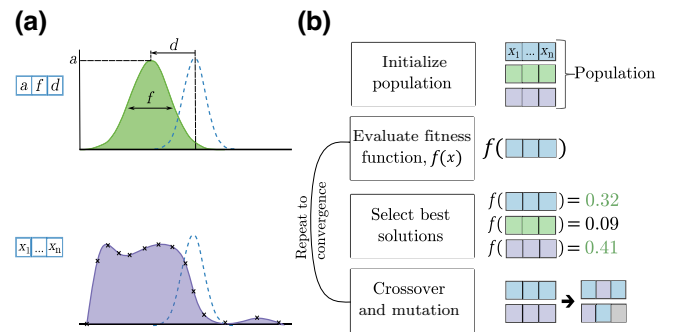


FIG. 2. (a) Different gene representations for different solutions. Only three genes are used to encode the Gaussian; the amplitude, a , the FWHM, f , and the delay with respect to the signal pulse, d . Sixteen evenly distributed points, x_1, \dots, x_{16} , encode the free-form pulse. The signal pulse is shown using the dashed blue line. (b) An overview of a genetic algorithm.

improvement in efficiency due to the genetic algorithm from the improvement in efficiency due to the pulse encoding, we run the genetic algorithm for both Gaussian and free-form solution representations and compare the results. This corresponds to performing a constraint optimization such as the following:

$$\hat{\theta} = \arg \max_{\theta} (\eta_{\text{int}}(\theta)), \quad (2)$$

where

$$\theta = \begin{cases} [a, f, d], & \text{if encoding a Gaussian pulse,} \\ [x_1, \dots, x_{16}], & \text{if encoding a free-form pulse,} \end{cases}$$

and η_{int} is the efficiency of the memory experiment performed, with a write pulse generated from θ (see Eq. 1).

In the Gaussian encoding, each solution is represented by three parameters: the amplitude of the Gaussian, a , the FWHM of the Gaussian, f , and the delay of the pulse, d , with respect to the incoming signal pulse [see Fig. 2(a)]. The a is allowed to take 50 discrete values normalized to between 0 and 1, to reduce the granularity with which we search the feature space, to reduce the number of epochs needed for convergence. Similarly, d and f are limited in their possible range of values to between -60 and 0 ns and $(0$ and 80 ns), respectively, both increasing in 1-ns intervals.

To implement a free-form pulse optimization, we represent a solution by 16 evenly distributed points, which are interpolated using a smoothing cubic spline method, to form a continuous waveform. Each point is limited to a range of values between -0.2 and 1 , which are discretized to 50 different levels. When the points are interpolated, any negative values in the waveform are clipped to 0. This is to ensure that steep gradients can be learned, while not generating nonphysical negative pulse intensities.

For each of the two encodings, an initial population of 60 individuals is randomly generated. For each individual, we generate the electrical waveform encoded by the genes and modulate the waveform into the laser beam by altering the rf power driving the AOM. Then we perform a memory experiment, as detailed in Sec. II A, using this waveform. The memory experiments for all individuals in a generation are recorded by an oscilloscope in a single shot and the trace is divided into individual memory experiments for evaluation.

As mentioned in Sec. II A, the internal efficiency is typically measured by detuning the laser, to ensure that the signal pulse has the same losses as the retrieval pulse. However, to measure the signal far detuned for each solution requires significant time overhead, which is deemed unfeasible for this experiment. Consequently, we evaluate the solutions on a fitness function which is faster to evaluate and is linearly proportional to the internal efficiency.




FIG. 3. The convergence of the free-form pulse, learned for a 18-ns-width signal. The fitness (light blue) and the related internal efficiency (dark blue) show the saturation of the learning process at generation 30. The pulse shapes for the best-performing pulse at each iteration are shown in purple; thus the final pulse is the end result. The green shows the number of new solutions tried per generation.

Specifically, to measure the fitness of each individual, we integrate the retrieved signal and normalize by a part of the input signal measured on a photodiode (PD2), before the cesium cell. The part of the input signal that is measured before the memory is integrated and is taken as a normalizing factor.

We select the solutions with the highest fitness to form the parents of the next generation of solutions. This is done by tournament selection with a tournament size of ten, until ten parents are chosen. Once selected, the parents are crossed over uniformly and each of the genes of the resulting children is randomly mutated with probability $p = 0.3$. Fifty-five new children are generated in this manner and form the next generation; five of the best-performing solutions survive unchanged. This process is repeated for 50 (25) epochs for the free-form (Gaussian) pulses. As a single evaluation of a generation takes approximately 1.5 mins, the time taken to run the whole experiment is $1.5 \text{ mins} \times \text{number of generations}$. To ensure the timely termination of the algorithm, without cutting off the evolution too early, we fix the number of generations executed. Given the comparatively small feature space to be searched in the Gaussian case, we find the algorithm to converge before 25 generations. In Fig. 3, we show a typical convergence plot for a single run of the genetic algorithm with free-form gene encoding, which illustrates the convergence of the algorithm well before 50 generations. Solutions that have been already evaluated are saved in a dictionary and their value recalled if they are to be reevaluated, resulting in strictly monotonous convergence.

III. RESULTS AND DISCUSSION

A. Efficiency optimization

We first run the genetic algorithm optimizing storage efficiency for different-width Gaussian input pulses. Here, we scan the FWHM of the signal pulse from 3.8 ns to 43 ns ($\tau_{\text{FWHM}}\gamma = \{1.9, 21.5\}$, $\gamma \approx 500$ MHz) and, for each signal, learn first the optimized Gaussian pulse and

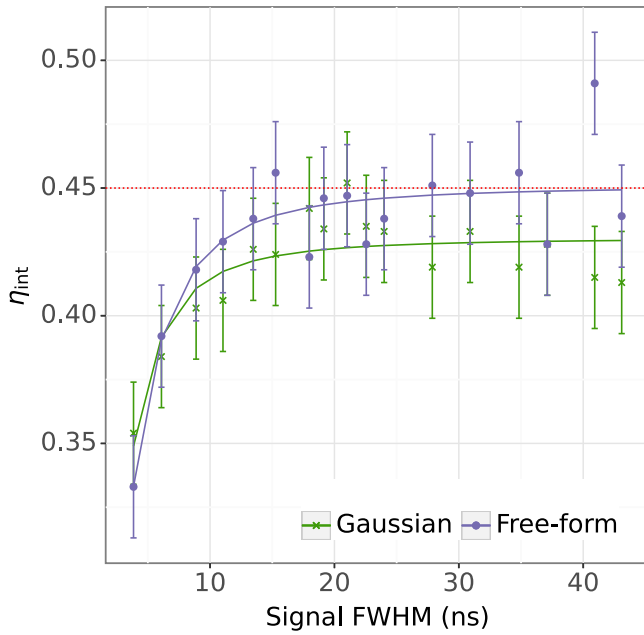


FIG. 4. The efficiencies of the learned pulses, for varying values of the signal FWHM. The free-form pulses (purple) look to perform on average, slightly better than the Gaussian pulse (green) but this improvement is not statistically significant. The fits are modeled as $\eta_{\text{int}} = \eta_{0,t_S} / \sqrt{1 + (4 \ln(2) / \Delta t_S \Delta \gamma)^2}$, where η_{0,t_S} is the maximal achievable efficiency and γ is the bandwidth of the memory (for the derivation, see Ref. [37]). The theoretically achievable efficiency at 200 ns is plotted in red, for a calculated OD = 12. We consider the measured 41-ns free-form efficiency to be an outlier with no physical meaning.

then the optimized free-form pulse. The results can be seen in Fig. 4. On average, the free-form pulse yields an improvement of 3(7)% over the learned Gaussian pulses. This extends the trend shown in Ref. [35], which shows that in a simulated three-level system, there is almost no difference in attainable efficiency between the free-form and Gaussian control pulses, for signal widths in the range $\tau_{\text{FWHM}}\gamma = \{0.001, 1.5\}$. The agreement between simulation and experiment is particularly surprising, given the presence of known physical effects, which are not accounted for in three-level systems. Effects such as the inhomogeneous broadening of the excited state, differences in the Rabi frequency, and the strength of coupling to different Zeeman levels are correlated with the pulse energy and thus the temporal shape of the pulses used in the memory. The larger the signal pulse is in duration, the more narrowband it becomes. As the EIT memory protocol is an adiabatic process, it is well suited to the storage of narrowband pulses—which can be seen in the saturation of the storage efficiency beyond 15-ns pulses. Conversely, the reduction of efficiency shown as the pulses become more broadband indicates the unsuitability of the memory protocol to the signal pulses to be stored, in agreement with

the literature [41]. We assess the measurement uncertainty by initially repeating a memory experiment using learned pulses and gauging the variance in efficiency, yielding a deviation of 0.015. Subsequently, we conduct efficiency-optimization procedures 10 times for both Gaussian and spline pulses, measuring deviations in the efficiency of the final learned pulses to be within the same order of magnitude; specifically, 0.015 and 0.016, respectively.

It seems as though the free-form pulses may yield only minimal improvements in the total storage efficiency, yet before accepting this conclusion one must also consider nuances of the pulse generation, particularly for smaller signal widths. Due to the 15-ns rise time of the AOM, for shorter signal pulses, it is not possible to modulate the optical field within the signal field for small signals. We are only able to modulate the field in the form of the falling flank of the AOM as it is switched off. Hence, it is not possible to definitively dismiss non-Gaussian pulses as nonoptimal for signal pulses smaller than $\tau_{\text{FWHM}}\gamma = 1.9$.

In Fig. 5, we show the final learned pulses for a subset of the signal widths. A general feature of all learned free-form pulses is a Gaussian-like falling-edge temporal overlap with the signal pulse. For signal widths above 18 ns, we observe a positively increasing trend in the distance between the peaks of the signal and control fields (see the insets to Fig. 5). Qualitatively, we see that the downward slope of the control pulse crosses first with the rising edge of the signal pulse, i.e., the optimal control pulse arrives before the signal. This behavior is also consistent for the Gaussian pulses and is consistent with the optimal pulse learned in Ref. [36].

To determine how important specific genes are to the efficiency, we use two indicators, the distribution of the state space explored during the learning process and the extent to which the genes of well-performing solutions vary. First, we consider the state-space search by plotting the value distribution of all the solutions (purple violin plot in Fig. 5). Genes where the importance of the value has little effect on the overall internal efficiency will have a wide value distribution, as seen in genes 2 and 7 of Fig. 5(a) and genes 2 and 4 of Fig. 5(b). Similarly, genes the value of which is important to the efficiency of the experiment will have a compact value distribution, such as most of the genes in the 11–16 range of genes. This is not a surprise, as high gene values would trigger early readout of the pulse, reducing the residual signal stored in the atoms and thus reducing the retrieval. Second, we calculate the variance for solutions that are at least 90% of the maximum fitness. These can be seen in blue in Fig. 5. We see a consistent trend across all learned pulses that the variance fluctuates in genes 1–9 and that beyond gene 10 the fluctuations reduce or are not present. This is an effect of preventing the aforementioned readout of the pulse. We see that the variance mirrors the distributions over the whole

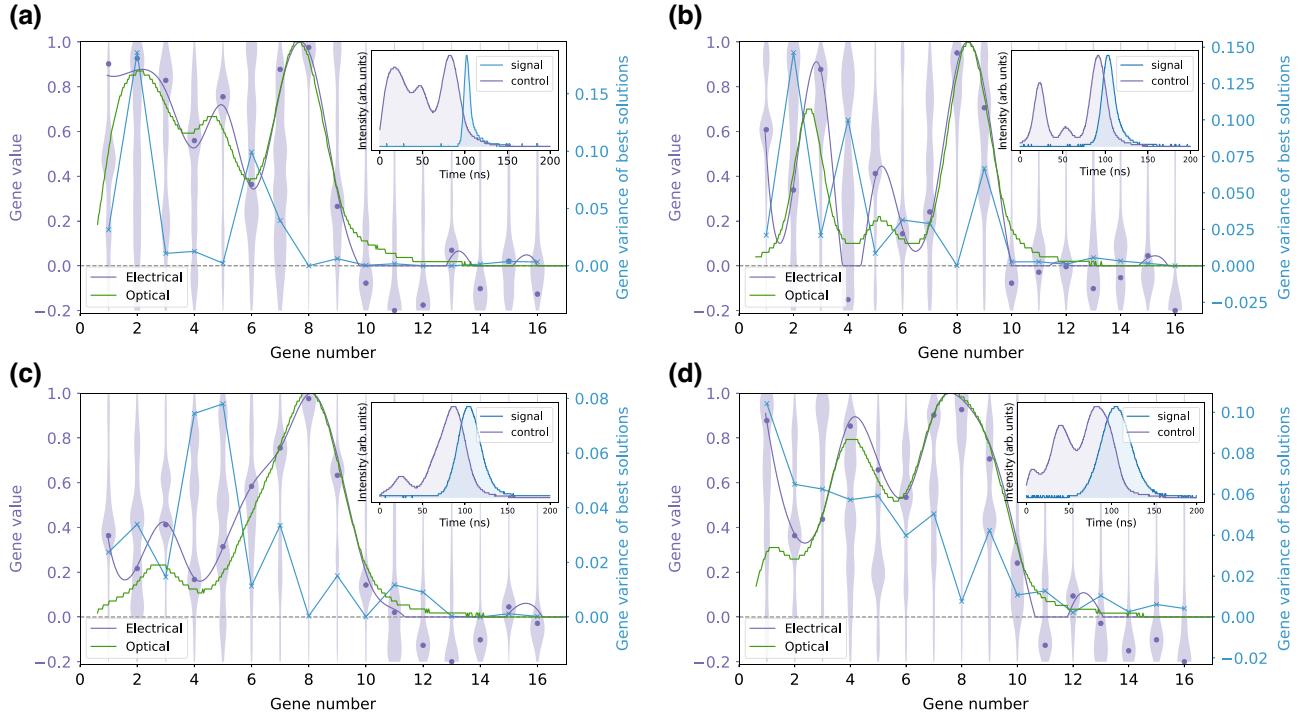


FIG. 5. The free-form pulse optimization. We show the best-learned pulses for a signal pulse with (a) 8.9-ns, (b) 18-ns, (c) 31-ns, and (d) 41-ns FWHM. The distribution of values explored for each gene throughout the whole optimization process is shown in a violin plot, with the final chosen gene values indicated as purple dots. The corresponding electrical waveform generated from the final gene values is shown in purple and its corresponding optical signal in green. The variance of solutions that have a fitness within 10% of the best-learned solution is plotted in blue. The lines joining the points have no physical significance and are intended as a guide to the eye. The temporal overlap of the signal and control pulses shows a consistent trend in the closing of the coherence window (inset).

searched space; i.e., genes that vary over the whole state space also have a high variance in the top 10% of solutions. Thus, one can conclude that solutions that have a low variance are important to the efficiency and thus should be the focus of further optimization efforts. We find the concentration of high-varying genes before the arrival of the signal pulse unsurprising, as these genes encode for time before the arrival of the signal pulse.

B. Energy optimization

One application that we demonstrate in this paper is to learn solutions that are limited in the total pulse energy.

Reference [37] shows the effect of the pulse energy on the signal-to-noise ratio of the retrieved pulses from the optical memory. By setting genes of the pulse that have no effect on the efficiency to high values, one increases the total pulse energy. This, in turn, inadvertently decreases the signal-to-noise ratio of the memory, with no efficiency payoff. To mitigate this effect in solutions learned by the genetic algorithm, one can reduce the allowed energy through an energy optimization. Indeed, a benefit to using a genetic algorithm is that we can set secondary optimization objectives, to learn optimal pulses given particularly restrained conditions.

To optimize the energy, we set the maximally learned pulse energy to be the area under the curve of the electrical pulses learned in the experiments described in Sec. III A. Then, we generate new solutions as described in Sec. II B but renormalize any pulses the electrical pulse area of which is larger than a specified limit, thereby setting a hard constraint on the allowed pulse energy of all the generated solutions. This leads to a modified constraint optimization function:

$$\tilde{\theta} = \arg \max_{\theta} (\eta'_{\text{int}}(\theta)), \quad (3)$$

where

$$\eta'_{\text{int}} = \begin{cases} \eta_{\text{int}}(\theta), & \text{if } I(\theta) \leq I(\hat{\theta})\alpha, \\ \eta_{\text{int}}(\theta/\beta), & \text{otherwise.} \end{cases}$$

Here, $I(\theta)$ is the integrated area under the curve of the pulse encoded by θ , $\hat{\theta}$ are the learned parameters from the efficiency optimization, α is a percentage factor of the maximum energy, $I(\hat{\theta})$, and the normalizing factor, β , is given by $\beta = I(\theta)/I(\hat{\theta})\alpha$. Here, we note that an advantageous aspect of employing a genetic algorithm in this experiment is its capacity to execute multiobjective optimization without requiring an explicit multiobjective loss

function. This circumvents the need for hyperparameter optimization aimed at guaranteeing that the discovered solutions are not Pareto dominant.

We choose the energy limit to be a progressively decreasing percentage of the learned pulse ($\alpha = \{0.3, 0.9\}$), as shown in Figs. 5(b) and 5(c). We carry out the energy optimization in the medium (18-ns) and long (31-ns) regimes and the results can be seen in Fig. 6. The upper pulse-energy limit corresponds to an optical pulse energy of $E_{G,18\text{ ns}} = 30.9\text{ nJ}$ and $E_{G,31\text{ ns}} = 33.0\text{ nJ}$ for the Gaussian pulses and $E_{F,18\text{ ns}} = 27.0\text{ nJ}$, $E_{F,31\text{ ns}} = 52.2\text{ nJ}$ for the free-form pulses.

We show that for a 31-ns signal pulse, we see a consistent trade-off in allowed pulse energy and attainable internal efficiency. Specifically, we see as the allowed pulse energy decreases, so too does the attainable optimized

efficiency. This suggests that without the need for an extra optimization objective, such as the power optimization presented in this section, the efficiency optimization alone already goes part of the way to learning the most power-efficient pulse. Moreover, when we consider the absolute energy of the learned pulses, we note that the maximum allowed energy for the free-form pulses is about 1.6 times larger than the corresponding maximum Gaussian energy. This means that, in absolute terms, the 60% energy free-form pulse and the 90% Gaussian pulse are limited to about the same energy and, accordingly, yield similar efficiencies. This supports the conclusion that for large signal sizes, Gaussians are good approximations of the most energetically efficient pulse. Furthermore, we see in Fig. 6(b) that the falling flank of both the learned free-form and Gaussian pulses remains consistent as the

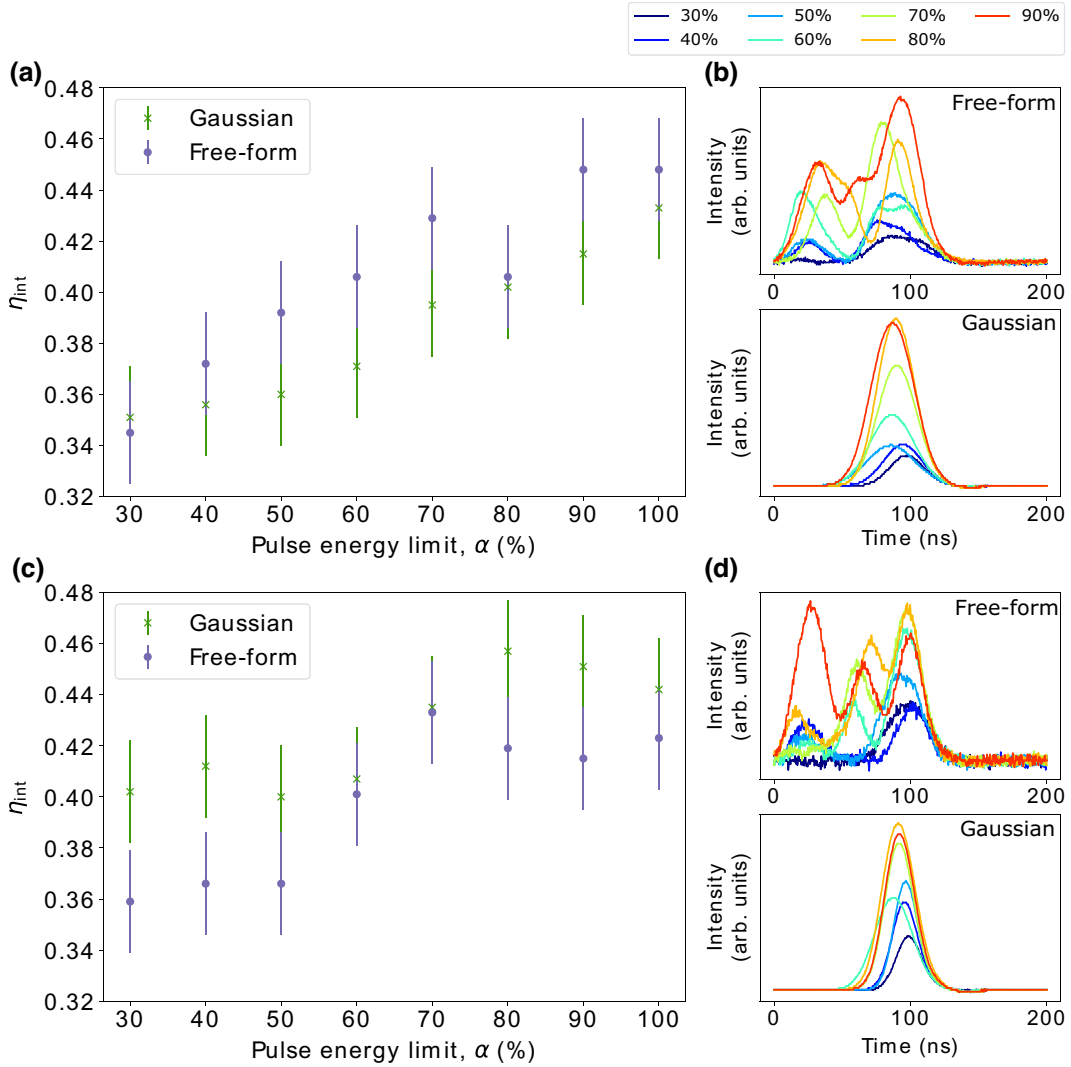


FIG. 6. The efficiencies of the energy-restricted pulses, for (a) 31-ns and (c) 18-ns FWHM signal pulses. The corresponding learned pulses are plotted for the free-form pulses [plotted in the upper panels of (b) and (d)] and Gaussian pulses [plotted in the lower panels of (b) and (d)], with a color gradient from blue to red (in energy terms, low to high).

energy is reduced, supporting the conclusions drawn from Sec. III A.

However, for a 18-ns signal pulse [see Fig. 6(c)], we are able to reduce the pulse energy of both the free-form and Gaussian pulses by 30% with a reduction of only 4(6)% from the efficiency at 100% pulse energy. For the free-form pulses, this reduction in energy corresponds to the reduction in the first part of the pulse (from 0 to 75 ns), while the falling flank remains largely constant. This supports our conclusions from the gene analysis that the first genes are not important factors in the efficiency. Indeed, we see that in the very-low-energy regime, we tend to a pulse shape that looks similar to the Gaussian. It is important to note that only Fig. 6(a) shows a signal pulse that is wide enough such that one would be able to modulate the control pulse within the signal pulse, due to limitations in the AOM rise time.

While this method is unable to reduce the pulse energy while maintaining efficiency for all signal widths, it does give us a methodology that allows us to specify a pulse energy, determined by the desired signal-to-noise ratio, and to learn optimal pulses for that energy level. This can be key in pushing the efficiency of warm quantum memories further.

IV. CONCLUSIONS

In summary, we show that a genetic algorithm can be used to optimize the write pulses of an optical quantum memory. The choice of a free-form pulse encoding gives an overall average improvement factor of 3(7)%, suggesting that Gaussian pulses are an acceptable approximation to optimal pulses. This agrees with the findings of Ref. [36]. Nonetheless, we demonstrate the merit of genetic algorithms in optimization with additional constraints, such as a total-pulse-energy limit. Here, we demonstrate that it is possible to reduce the pulse energy by 30% in some cases, without a large compromise in efficiency. This work focuses on the measurement of the write-control-pulse optimization for Gaussian-shaped signal pulse; however, other works investigating signal shapes closer to the emission of single-photon sources (see Ref. [30]) indicate the value of extending this approach to different signal shapes. Moreover, future work considering the learning and analysis of an optimal read pulse may be interesting to determine the validity of the time-reversal assumption, often used in theoretical modeling [27]. Finally, this is a platform-agnostic approach that can be applied to a wide range of atomic and molecular physics experiments, supporting the further development of a range of high-efficiency optical quantum memories.

ACKNOWLEDGMENTS

This work was funded by the German Ministry of Education and Research (BMBF) project “Quantum

Tokens Based on Alkali Metals and Xenon” (Q-ToRX) and the Deutsche Forschungsgemeinschaft through Grant No. 448532670. E.R. acknowledges funding through the Helmholtz Einstein International Berlin Research School in Data Science (HEIBRiDS).

APPENDIX A: GENETIC ALGORITHM DETAILS

In this appendix, we elaborate on the processes used to select the parents and generate the children of the next generation. Once the fitnesses of each solution have been evaluated, one must consider a method of choosing which solutions will be selected as parents. There are many different selection methods available; in this work, we chose tournament selection, as it is easy to conceptualize and select the selection pressure (the likelihood that suboptimal solutions are selected). In tournament selection, a subsection of the solutions are drawn at random. The highest-performing solution from the subset is selected as a parent for the next generation. This process is repeated until one has generated a number of parents, specified by the “number of parents mating” hyperparameter. The list of parents is traversed sequentially and each pair of parents is used to generate a solution of the next generation by crossover, i.e., first parents 1 and 2 are selected, then parents 2, 3, etc. Once two parents have been selected and their genes are crossed over, there are several methods for performing crossover, which can be selected depending on the physical meaning of the genes. We select a uniform crossover, such that for each gene, one of the genes of the two parents is chosen with equal probability and that gene is copied across to the other. Once crossed over, each gene of the child is mutated with a probability p and the resulting solution is taken as one element of the population in the next generation. Table I lists the hyperparameters that are used in the experiment. The hyperparameters are selected either on a trial-and-error basis or, in the case of solutions per population and generations, on the basis of on system limitations such as the AFG waveform memory.

TABLE I. The hyperparameters used to run the genetic algorithm [40]. The values quoted in parenthesis are the parameters used for the Gaussian optimization.

Hyperparameter	Free-form (Gauss)
Number of genes	16 (3)
Number of generations	50 (25)
Number of solutions per population	60
Number of parents mating	10
Selection type	Tournament
Tournament size	10
Elitism size	5
Crossover type	Uniform
Mutation type	Random
Mutation probability	0.3

- [1] J. Nunn, K. Reim, K. C. Lee, V. O. Lorenz, B. J. Sussman, I. A. Walmsley, and D. Jaksch, Multimode memories in atomic ensembles, *Phys. Rev. Lett.* **101**, 260502 (2008).
- [2] N. Sangouard, C. Simon, H. de Riedmatten, and N. Gisin, Quantum repeaters based on atomic ensembles and linear optics, *Rev. Mod. Phys.* **83**, 33 (2011).
- [3] M. Gündoğan, J. S. Sidhu, V. Henderson, L. Mazzarella, J. Wolters, D. K. L. Oi, and M. Krutzik, Proposal for space-borne quantum memories for global quantum networking, *npj Quantum Inf.* **7**, 128 (2021).
- [4] J. Wallnöfer, F. Hahn, F. Wiesner, N. Walk, and J. Eisert, Faithfully simulating near-term quantum repeaters, *PRX Quantum* **5**, 010351 (2024).
- [5] J.-M. Mol, L. Esguerra, M. Meister, D. E. Bruschi, A. W. Schell, J. Wolters, and L. Wörner, Quantum memories for fundamental science in space, *Quantum Sci. Technol.* **8**, 024006 (2023).
- [6] A. D. Manukhova, K. S. Tikhonov, T. Y. Golubeva, and Y. M. Golubev, Noiseless signal shaping and cluster-state generation with a quantum memory cell, *Phys. Rev. A* **96**, 023851 (2017).
- [7] M. Gündoğan, *et al.*, Topical white paper: A case for quantum memories in space, [arXiv:2111.09595](https://arxiv.org/abs/2111.09595).
- [8] L. Jaurigue, E. Robertson, J. Wolters, and K. Lüdge, Reservoir computing with delayed input for fast and easy optimisation, *Entropy* **23**, 1560 (2021).
- [9] Y. Lei, F. Kimiaee Asadi, T. Zhong, A. Kuzmich, C. Simon, and M. Hosseini, Quantum optical memory for entanglement distribution, *Optica* **10**, 1511 (2023).
- [10] H. P. Specht, C. Nölleke, A. Reiserer, M. Uphoff, E. Figueroa, S. Ritter, and G. Rempe, A single-atom quantum memory, *Nature* **473**, 190 (2011).
- [11] M. Gündoğan, P. M. Ledingham, K. Kutluer, M. Mazzer, and H. de Riedmatten, Solid state spin-wave quantum memory for time-bin qubits, *Phys. Rev. Lett.* **114**, 230501 (2015).
- [12] T.-S. Yang, Z.-Q. Zhou, Y.-L. Hua, X. Liu, Z.-F. Li, P.-Y. Li, Y. Ma, C. Liu, P.-J. Liang, X. Li, Y.-X. Xiao, J. Hu, C.-F. Li, and G.-C. Guo, Multiplexed storage and real-time manipulation based on a multiple degree-of-freedom quantum memory, *Nat. Commun.* **9**, 3407 (2018).
- [13] Y.-F. Pu, N. Jiang, W. Chang, H.-X. Yang, C. Li, and L.-M. Duan, Experimental realization of a multiplexed quantum memory with 225 individually accessible memory cells, *Nat. Commun.* **8**, 15359 (2017).
- [14] S.-J. Yang, X.-J. Wang, X.-H. Bao, and J.-W. Pan, An efficient quantum light-matter interface with sub-second lifetime, *Nat. Photonics* **10**, 381 (2016).
- [15] X.-H. Bao, A. Reingruber, P. Dietrich, J. Rui, A. Dück, T. Strassel, L. Li, N.-L. Liu, B. Zhao, and J.-W. Pan, Efficient and long-lived quantum memory with cold atoms inside a ring cavity, *Nat. Phys.* **8**, 517 (2012).
- [16] Y.-H. Chen, M.-J. Lee, I.-C. Wang, S. Du, Y.-F. Chen, Y.-C. Chen, and I. A. Yu, Coherent optical memory with high storage efficiency and large fractional delay, *Phys. Rev. Lett.* **110**, 083601 (2013).
- [17] Y.-W. Cho, G. T. Campbell, J. L. Everett, J. Bernu, D. B. Higginbottom, M. T. Cao, J. Geng, N. P. Robins, P. K. Lam, and B. C. Buchler, Highly efficient optical quantum memory with long coherence time in cold atoms, *Optica* **3**, 100 (2016).
- [18] P. Vernaz-Gris, K. Huang, M. Cao, A. S. Sheremet, and J. Laurat, Highly-efficient quantum memory for polarization qubits in a spatially-multiplexed cold atomic ensemble, *Nat. Commun.* **9**, 363 (2018).
- [19] Y.-F. Hsiao, P.-J. Tsai, H.-S. Chen, S.-X. Lin, C.-C. Hung, C.-H. Lee, Y.-H. Chen, Y.-F. Chen, I. A. Yu, and Y.-C. Chen, Highly efficient coherent optical memory based on electromagnetically induced transparency, *Phys. Rev. Lett.* **120**, 183602 (2018).
- [20] Y. Wang, J. Li, S. Zhang, K. Su, Y. Zhou, K. Liao, S. Du, H. Yan, and S.-L. Zhu, Efficient quantum memory for single-photon polarization qubits, *Nat. Photonics* **13**, 346 (2019).
- [21] M. Sabooni, Q. Li, S. Kröll, and L. Rippe, Efficient quantum memory using a weakly absorbing sample, *Phys. Rev. Lett.* **110**, 133604 (2013).
- [22] M. P. Hedges, J. J. Longdell, Y. Li, and M. J. Sellars, Efficient quantum memory for light, *Nature* **465**, 1052 (2010).
- [23] L. Meßner, E. Robertson, L. Esguerra, K. Lüdge, and J. Wolters, Multiplexed random-access optical memory in warm cesium vapor, *Opt. Express* **31**, 10150 (2023).
- [24] A. V. Gorshkov, T. Calarco, M. D. Lukin, and A. S. Sørensen, Photon storage in Λ -type optically dense atomic media. IV. Optimal control using gradient ascent, *Phys. Rev. A* **77**, 043806 (2008).
- [25] J. Guo, X. Feng, P. Yang, Z. Yu, L. Q. Chen, C.-H. Yuan, and W. Zhang, High-performance raman quantum memory with optimal control in room temperature atoms, *Nat. Commun.* **10**, 148 (2019).
- [26] J. Wolters, G. Buser, A. Horsley, L. Béguin, A. Jöckel, J.-P. Jahn, R. J. Warburton, and P. Treutlein, Simple atomic quantum memory suitable for semiconductor quantum dot single photons, *Phys. Rev. Lett.* **119**, 060502 (2017).
- [27] A. V. Gorshkov, A. André, M. D. Lukin, and A. S. Sørensen, Photon storage in Λ -type optically dense atomic media. II. Free-space model, *Phys. Rev. A* **76**, 033805 (2007).
- [28] I. Novikova, N. B. Phillips, and A. V. Gorshkov, Optimal light storage with full pulse-shape control, *Phys. Rev. A* **78**, 021802(R) (2008).
- [29] N. B. Phillips, A. V. Gorshkov, and I. Novikova, Optimal light storage in atomic vapor, *Phys. Rev. A* **78**, 023801 (2008).
- [30] M. T. Rakher, R. J. Warburton, and P. Treutlein, Prospects for storage and retrieval of a quantum-dot single photon in an ultracold ^{87}Rb ensemble, *Phys. Rev. A* **88**, 053834 (2013).
- [31] T. Hornung, R. Meier, D. Zeidler, K.-L. Kompa, D. Proch, and M. Motzkus, Optimal control of one- and two-photon transitions with shaped femtosecond pulses and feedback, *Appl. Phys. B: Lasers Opt.* **71**, 277 (2000).
- [32] D. Zeidler, S. Frey, K.-L. Kompa, and M. Motzkus, Evolutionary algorithms and their application to optimal control studies, *Phys. Rev. A* **64**, 023420 (2001).
- [33] R. S. Judson and H. Rabitz, Teaching lasers to control molecules, *Phys. Rev. Lett.* **68**, 1500 (1992).
- [34] V. C. Gregoric, X. Kang, Z. C. Liu, Z. A. Rowley, T. J. Carroll, and M. W. Noel, Quantum control via a genetic algorithm of the field ionization pathway of a Rydberg electron, *Phys. Rev. A* **96**, 023403 (2017).

- [35] K. Shinbrough, B. D. Hunt, and V. O. Lorenz, Optimization of broadband Λ -type quantum memory using gaussian pulses, *Phys. Rev. A* **103**, 062418 (2021).
- [36] K. Shinbrough and V. O. Lorenz, Variance-based sensitivity analysis of Λ -type quantum memory, *Phys. Rev. A* **107**, 033703 (2023).
- [37] L. Esguerra, L. Meßner, E. Robertson, N. V. Ewald, M. Gündoğan, and J. Wolters, Optimization and readout-noise analysis of a warm-vapor electromagnetically-induced-transparency memory on the Cs D_1 line, *Phys. Rev. A* **107**, 042607 (2023).
- [38] J. H. Holland, *Adaptation in Natural and Artificial Systems—An Introductory Analysis with Applications to Biology, Control, and Artificial I* (The MIT Press, Cambridge, Massachusetts, 1975).
- [39] M. Mitchell, *An Introduction to Genetic Algorithms*, 1st ed., A Bradford book (The MIT Press, Cambridge, Massachusetts, 1998).
- [40] A. F. Gad, PyGAD: An intuitive genetic algorithm PYTHON library, *Multimed. Tools Appl.* **83**, 58029 (2024).
- [41] A. Rastogi, E. Saglamyurek, T. Hrushevskiy, S. Hubele, and L. J. LeBlanc, Discerning quantum memories based on electromagnetically-induced-transparency and Autler-Townes-splitting protocols, *Phys. Rev. A* **100**, 012314 (2019).