



Single-shot time-folded fluorescence lifetime imaging

Valentin Kapitany^{a,1,2} , Vytautas Zickus^{a,b,1} , Areeba Fatima^a, Guillem Carles^a , and Daniele Faccio^{a,2}

Edited by John T. Fourkas, University of Maryland, College Park, MD; received August 26, 2022; accepted March 8, 2023, by Editorial Board Member John A. Rogers

Fluorescence lifetime imaging is an important tool in bioimaging that allows one to detect subtle changes in cell dynamics and their environment. Most time-domain approaches currently involve scanning a single illumination point across the sample, which can make imaging dynamic scenes challenging, while single-shot “rapid lifetime determination” can suffer from large uncertainties when the lifetime is not appropriately sampled. Here, we propose a time-folded fluorescence lifetime imaging microscopy (TFFLIM) approach, whereby a time-folding cavity provides multiple spatially sheared replicas of the lifetime, each shifted temporally with respect to a fixed time gate. This provides a robust, single-shot FLIM approach that we experimentally validate across a broad lifetime range on fluorescent beads and *Convallaria* samples.

FLIM | single-shot | time-folded cavity | physics-inspired neural network

Fluorescence lifetime imaging quantifies the time-dependent properties of fluorophores (1–3) and is typically employed for imaging biological phenomena. FLIM measurements are utilized to extract information about the local environment of the fluorophores, such as concentration of oxygen, pH, as well as to reveal protein–protein interactions, changes in membrane tension, etc. (3–7). Despite the fact that FLIM has been recently shown to benefit clinical diagnostics (8–10), wider adoption of the technique in clinical settings is still hindered due to limited field of view (FOV) and slow acquisition speeds of commercial FLIM systems (10). The main bottleneck in conventional systems is the point-scanning image acquisition, which prohibits the acquisition of instantaneous full FOV information. To tackle this, recent innovation has focused on improving the acquisition speeds of confocal scanning systems, in particular with resonant scanning and spinning-disc systems (11) in conjunction with detection approaches based, for example, on photomultiplier tubes (PMT) with time-correlated single photon counting electronics (TCSPC). Nonetheless, if dynamic scenes and large FOVs are desired, or plane-illumination schemes like light-sheet imaging or total internal reflection fluorescence (TIRF) microscopy are used, wide-field systems would be the solution.

Wide-field single-photon avalanche diode (SPAD) arrays have gained popularity both in TCSPC (12) and time-gated operation (13, 14). Wide-field FLIM systems have also been shown that utilize microchannel-plate gated optical intensifiers in addition to arrayed PMT detectors (15) or charge-coupled device (CCD) cameras (16, 17). Such systems benefit from a high-fill factor and low noise and were designed for low-light imaging applications. The fluorescence decay information in such systems is extracted by scanning a gate that samples different temporal sections of the decay. In the extreme case where only 2 gates are used, the acquired data lends itself to relatively fast and simple lifetime recovery techniques referred to as Rapid Lifetime Determination (RLD) (18).

Recently, variations of RLD have been proposed that bypass scanning a gate, including a novel wide-field FLIM system utilizing Pockels cells and a 4f optical cavity (19), which creates a time-integrated measurement in a single snapshot image. The cost of this approach is that 16 replicas of the FOV are imaged onto the detector, therefore potentially limiting total FOV or resolution.

Attempts have been made at using a single, gated camera in combination with temporal streaking of the fluorescent signal in order to determine its lifetime. For example, combining compressed ultrafast photography (CUP) (20) with a dual camera detection scheme has been demonstrated to allow wide-field streaked imaging (21). However, CUP and related approaches rely on streak cameras and leave the need for setups with simpler or cheaper components. Heshmat et al. have also shown that it is possible to retrieve the temporal evolution of systems that evolve at the picosecond-to-nanosecond scale with a time-folding approach whereby multiple time-delayed discrete replicas of a laser pulse result from its periodic optical round trips in an optical cavity (22).

Significance

This paper describes a single-shot fluorescence lifetime imaging (FLIM) method. We use an optical cavity to create temporally delayed and spatially sheared replicas of the fluorescent decay signal onto a time-gated intensified charged-coupled device (iCCD). This modality allows different portions of the decay signal to be sampled in parallel by a single camera, bypassing the convention of scanning a gate. The technique is validated experimentally on fluorescent beads and cells.

Author affiliations: ^aExtreme Light, School of Physics & Astronomy, University of Glasgow, Glasgow G12 8QQ, UK; and ^bPlasmonics and Nanophotonics Laboratory, Department of Laser Technologies, Center for Physical Sciences and Technology, Vilnius LT-10257, Lithuania

Author contributions: G.C. and D.F. designed research; V.Z. performed experiments; V.K. and A.F. contributed new reagents/analytic tools; V.K., A.F., and G.C. analyzed data; and V.K. and V.Z. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission. J.T.F. is a guest editor invited by the Editorial Board.

Copyright © 2023 the Author(s). Published by PNAS. This open access article is distributed under Creative Commons Attribution License 4.0 (CC BY).

¹V.K. and V.Z. contributed equally to this work.

²To whom correspondence may be addressed. Email: valentin.kapitany@glasgow.ac.uk or Daniele.Faccio@glasgow.ac.uk.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2214617120/-/DCSupplemental>.

Published April 12, 2023.

We propose a FLIM that sits in between the discrete sampling of RLD and the continuous streaking of CUP-based approaches. We use a purely optical, time-folded cavity to shear a fluorescent signal in space and time and replicate it within our field of view. This approach retains most of the field of view, while also producing multiple, discrete time-gated samples of the fluorescent signal. These samples act as a series of parallel measurements, similar to taking multiple RLD acquisitions of the same sample with different gate positions at the same time.

Lifetime images are then retrieved by an inverse retrieval scheme and a physics-inspired neural network, for predicting lifetime robustly and quickly, respectively. This approach is validated on various samples over a broad range of lifetimes.

Results

Time-Folded Cavity. The proposed experimental layout is shown in Fig. 1A. A uniform sample of fluorescent molecules will decay according to an exponential probability density function. Our optical cavity replicates this decay signal and retards each replica in time by ~ 0.59 ns, as shown in Fig. 1B. A tilted mirror M' in the cavity adds an angular offset to the signal as it leaves the cavity. The replicas pass through a lens; hence, this angular offset becomes a lateral translation on the iCCD. Consequently, the iCCD intensifies different portions of the signal on different pixels.

The noise-free time-folded lifetime signal at a pixel $s_{i,j}$ on the iCCD is

$$s_{i,j} = r_0 \int G(t) A_{i,j} e^{\frac{t_0-t}{\tau_{i,j}}} H(t-t_0) dt + r_1 \int G(t) A_{i,j-1y} e^{\frac{t_0+t_c-t}{\tau_{i,j-y}}} H(t-t_0-t_c) dt + \dots + r_N \int G(t) A_{i,j-Ny} e^{\frac{t_0+Nt_c-t}{\tau_{i,j-Ny}}} H(t-t_0-Nt_c) dt, \quad [1]$$

where r_n for $n \in 0, 1, \dots, N$ is a constant, which represents the power transmitted in the n -th round trip through the beamsplitter. The splitting ratio at the excitation wavelength was experimentally measured as 51.5:48.5. $G(t)$ is the temporal profile of the time-gated intensifier on the iCCD camera. A is the amplitude of the decay, and τ is the corresponding lifetime. $H(t-\alpha)$ is the Heaviside step function. The N terms represent N round trips; we determined $N = 5$ to be a good cutoff point,

i.e., further round trips had a poor signal-to-noise ratio (SNR) or were not visible. The image shear, measured in pixels per round trip, is y . At each round trip, the replicas are also delayed by the cavity round trip time, t_c .

In parallel to the iCCD acquisition, a CMOS camera also acquires an image where each CMOS pixel i, j records an intensity

$$q_{i,j} = k \int A_{i,j} e^{\frac{t_0-t}{\tau_{i,j}}} dt, \quad [2]$$

which corresponds to the time-integrated fluorescence recorded at each pixel. Here, k is a factor that represents the amplitude ratio of the measurement obtained on the iCCD and CMOS cameras. This arises due to the product of the splitting ratio of BS_1 , the pixel sizes, and fill factors of the two cameras, and their respective quantum yields at the given emission wavelength. As such, it is constant over our experiments.

Uncertainty Analysis. We assess the viability of our imaging modality by analyzing the lower bound of our prediction uncertainty. This lower bound is found by considering a Bayesian model and by assuming that the noise in the sCMOS is negligible compared to the iCCD. We then calculate the probability of a true sample lifetime τ given a lifetime estimate $\hat{\tau}$ obtained from a set of noisy iCCD measurements \hat{s} (i.e., $\hat{s}(\tau)$):

$$p(\tau|\hat{\tau}) \equiv p(\tau|\hat{s}) = \frac{p(\hat{s}|\tau)p(\tau)}{p(\hat{s})}. \quad [3]$$

This formalism allows us to derive lifetime uncertainty from the well-studied noise properties of an iCCD; *Methods* for details.

We evaluate the uncertainty lower bound of our system for four different example photon budgets in Fig. 2A. With 5,000 photons (incident on the iCCD), our lifetime prediction error is less than 10% for lifetimes between 0.36 and 2.6 ns. For a photon budget of 50,000, the same is true for over 2 orders of magnitude of lifetimes, from 0.1 to 10 ns. In the presence of thermal noise or added Poissonian noise on the CMOS, we expect the same overall trend in lifetime uncertainties, at slightly higher values.

Data Analysis. We propose two methods to retrieve the lifetime map. One is inverse retrieval (IR) with gradient descent, which is computationally slow but tractable. The other is a physics-inspired convolutional neural network (CNN) trained on simulated data; this method is fast yet intractable.

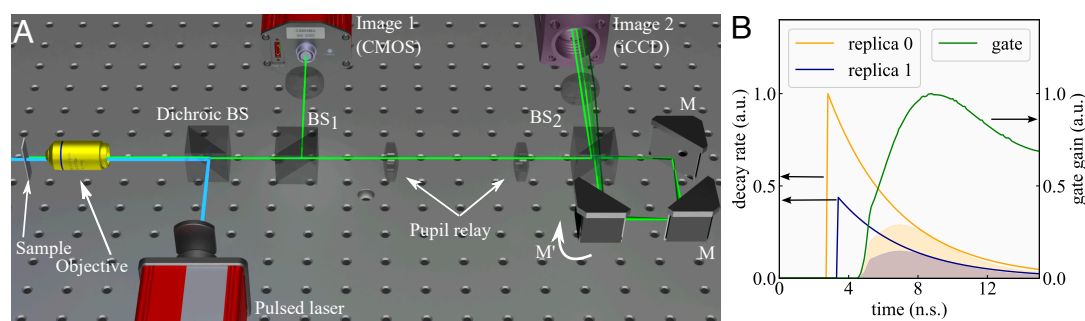


Fig. 1. (A) The system consists of an objective with epifluorescence illumination, 2 cameras (standard CMOS, and a gated iCCD), and a time-folding optical cavity assembled from a beamsplitter (BS') and 3 mirrors, 1 of which (M') controls the direction of the replicas and their separation in the image plane. A 1:1 telescope relays the pupil of the objective into the cavity to minimize vignetting and cropping of the FOV. (B) The optical cavity creates replicas of the original signal with an added time lag, caused by the cavity's round trip time. Each replica is translated spatially by M' , and its amplitude shrinks according to the splitting ratio of BS_2 . The iCCD signal measured from each replica is represented by the area of the shaded region under the curves; this area is the dot product of the exponential decay and the gate, shown truncated here for visualization; *SI Appendix* for the full gate shape.

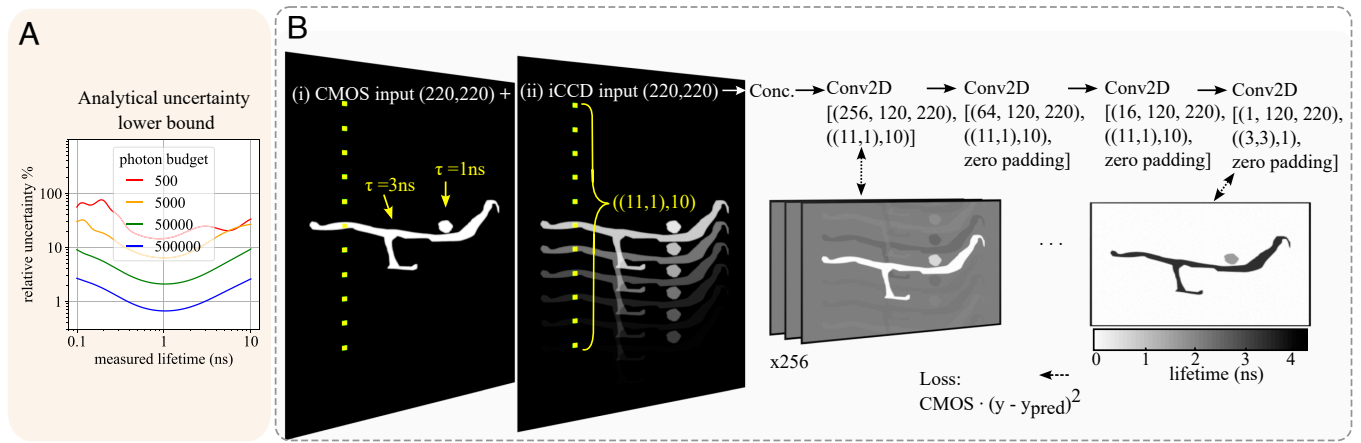


Fig. 2. (A) We show the lower bound on the relative uncertainty of the lifetime prediction (calculated using Eq. 3) for photon budgets of 500, 5,000, 50k, and 500k photons (incident on the iCCD). Similarly to RLD (18), our measurements have minimal uncertainty around some lifetime domain, which can be controlled by the position of the gate with respect to the excitation pulse, and the time delay between each sheared image. Our system is optimized to measure ≈ 1 ns signals. (B) Outline of the CNN. (i and ii) The CMOS and iCCD images are shown. They have two lifetime populations: a small sphere with a 1-ns lifetime and a long, thin structure with a 3-ns lifetime. The lifetime of the sample affects the relative intensities of subsequent replicas on the iCCD, with longer lifetime samples losing more intensity between one replica and the next. The dilated Conv2D kernel is shown in yellow; a dilation rate of 10 means that only every 10th kernel value is nonzero, creating a sparse receptive field. The inputs are concatenated, then processed by 3 Conv2D layers with kernel sizes of (11×1) and dilation rates of 10. The first hidden layer is shown; it consists of 256 feature maps. Since the inputs are not zero-padded, the size of the feature map decreases compared to the inputs. The output is a lifetime prediction. Training loss is mean squared error weighted by the corresponding intensity values.

The IR pipeline optimizes our guess of the lifetime map by iteratively computing the expected iCCD signal for a given lifetime map and calculating the gradient of mean squared error between this expected iCCD signal and the measurement with respect to the lifetime (called the data fidelity term), with gradient descent to update the lifetime guess. This process is analytically tractable. It can however stagnate in local minima of the mean squared error landscape. Therefore, we constrain the IR via an L2-norm regularizer on the lifetime map. Our optimization problem searches for the optimal solution $\hat{\tau}$:

$$\hat{\tau} = \arg \min_{\tau} C(\tau), \text{ where} \quad [4]$$

$$C(\tau) = \|\mathbf{P}(\tau) - \hat{s}\|_2 + \alpha \|\tau\|_2^2$$

subject to $\tau \geq 0$.

The first term of the cost function $C(\tau)$ represents the data fidelity term, and the second term represents the regularizer. The inverse retrieval is thus a nonlinear ridge regression problem.

Our CNN algorithm, shown in Fig. 2B, is trained on a set of synthetic iCCD and CMOS image pairs. It uses 2D convolutional (Conv2D) layers with dilated kernels. Dilation ensures explicit causal dependence of the feature maps on the inputs: the sparse receptive field of the kernel means information is drawn only from pixels with causal dependence in the forward model, minimizing overfitting and matching the physics of the problem. Since the convolutional kernels maintain this dilated shape until the penultimate layer, this causal dependence persists until the deeper layers. The final Conv2D layer's (3×3) kernels mimic sliding window binning, commonly used in lifetime fitting to increase the SNR. Training lifetime labels are in the range of 0.1 to 8 ns. For more details on the IR and the CNN, *Methods*.

We apply both IR and CNN to predict the lifetimes of the center of the field of view. In this region, we have information about all the replicas that enter and leave the field of view, allowing us to predict lifetimes with minimal uncertainty.

Validation on Beads. We prepared samples of fluorescent beads of 2 μm and 4 μm diameters with corresponding expected lifetimes

of 2.1 ns and 4 ns measured on a separate FLIM system (*Methods* for details). We show results for an FOV containing several such beads in Fig. 3. Exposure times were 1 s.

The 2- and 4- μm beads show bimodal lifetime distributions and are clearly distinguishable by their lifetimes. Moreover, both the neural network and IR yield very similar lifetime distributions, cross-validating one another. We fitted bimodal skewed Gaussian distributions to intensity-weighted (the weighting is used to remove dark pixels) histograms of the pixel values. For the CNN, peak values of 2.0 and 4.2 ns with standard deviations of 0.1 and 0.7 ns were obtained, for the small and large beads, respectively. IR yielded values of 2.0 ± 0.0 and




Fig. 3. (A and B) We show a mixture of 2 and 4 μm beads on the CMOS and iCCD. (C and D) The CNN predicts the two lifetime modes present in the sample clearly. The predicted lifetime distribution is fitted with a bimodal skewed Gaussian, yielding populations of 2.0 ± 0.1 and 4.2 ± 0.7 ns. (E and F) The IR predicts similar lifetimes, distributed as 2.0 ± 0.0 and 3.6 ± 0.7 ns.




Fig. 4. Lifetime map of *Convallaria* stained with acridine orange. The CNN and inverse retrieval show remarkably similar lifetime patterns, cross-validating one another. The CNN output is a little smoother, while the inverse retrieval has sharper edges but also more granularity; this can be attributed to the final layer of the CNN acting as a form of (learned) sliding window averaging. Reconstructions were intensity-weighted, as the regions in between cell walls expressed low signal. (A), (i and ii) The CMOS and iCCD images of the first region of interest. (iii and iv) The global lifetime distributions were 1.30 ± 0.09 and 1.29 ± 0.11 ns for the CNN and IR, respectively. (B), (i and ii) CMOS and iCCD images of the second region of interest. (iii and iv) In the second region of interest, the global lifetime distributions were 1.34 ± 0.11 and 1.30 ± 0.15 ns for the CNN and IR, respectively. We note that the artifacts produced by the inverse retrieval are largely suppressed by the CNN's final layer, but sharp edges and high-frequency lifetime variations are also smoothed.

3.6 ± 0.7 ns, respectively. These values are in good agreement with the expected lifetimes. The prediction uncertainties of the longer lifetime beads are greater than that of the shorter lifetime ones, as expected from Fig. 2A.

Testing on Cells. We then tested our model on *Convallaria* cells stained with acridine orange. Mean lifetimes of 1.30 and 1.29 ns for region-of-interest (ROI) 1 (Fig. 4A) were obtained with the CNN and IR, respectively. In ROI 2 (Fig. 4B), the mean lifetimes predicted by the CNN and IR were 1.34 and 1.30 ns, respectively. These distributions closely match our previously published results of 1.29 ± 0.49 ns (13) on the same sample type. Moreover, the samples do not have a flat lifetime, and the CNN and IR agree on the local lifetime pattern, cross-validating one another. This demonstrates that even though these samples are not spatially sparse like the beads (thus, their time-folded replicas on the iCCD overlap), TFFLIM is capable of predicting lifetimes with good fidelity.

Discussion

We have demonstrated a method for obtaining fluorescence lifetime images using a time-folded optical cavity in conjunction with an iCCD camera and a CMOS camera. We derive an inverse modeling approach and designed a physics-inspired neural network architecture, thus predicting sample lifetimes for this imaging scheme. TFFLIM samples the signal for different delays, mimicking parallel time-gated measurements, which are acquired in a single snapshot. We validated our approach on beads and tested it on biological samples, obtaining lifetimes in the expected ranges for both. We mathematically demonstrate that the cavity delivers a relatively uniform prediction uncertainty for a wide lifetime range.

The simultaneous spatial and temporal shearing implies that all the information required for retrieving the final lifetime image is acquired in a single measurement. This yields an advantage over scanning time-gated approaches, which sample the signal sequentially for various gate positions. Finally, since we acquire multiple time-gated measurements with different gate positions

in parallel, our method could theoretically be expanded to single-shot multiexponential retrieval, analogously to multiexponential RLD (23).

Materials and Methods

Setup. The cavity was assembled from a 50:50 beamsplitter cube (Thorlabs CCM1-BS013/M), two right-angle turning prism mirrors (Thorlabs CCM1-E02/M), an elliptical mirror mounted with a length-wise extended 3D printed mirror mount (Thorlabs H45E1) on a kinematic mount (Thorlabs KM100CP/M) allowing for fine adjustment of replica separation and angle. The imaging optics consisted of a pair of $f = 250$ -mm achromatic doublets (Thorlabs AC254-250-A-ML) which relayed the pupil of a Zeiss 40×0.75 NA microscope objective. Replicas were imaged onto an iCCD (Andor iStar 334T) with an $f = 500$ -mm achromatic doublet (Thorlabs AC254-500-A-ML).

The replica separation and direction can be found experimentally from the 2D Fourier transform of the iCCD image. The shear direction (and tilt from the desired direction) of the replicas is important as the reconstruction algorithm implicitly assumes that this direction is known. Small deviations from the expected shear direction (even of a few degrees) can lead to artifacts in the form of bands of higher lifetimes that can become very evident for larger angular deviations. The 2D Fourier transform allows us to precisely measure tilt in the shear direction, as well as the replica separation, and subsequently allows the iCCD image to be corrected before processing. In *SI Appendix*, we show examples of sheared images with and without a tilt alongside the relative reconstructions.

A 30:70 beamsplitter cube (Thorlabs BS019) before the pupil relay provided the imaging path for the CMOS camera (Thorlabs CS895MU) with $f = 200$ -mm achromatic doublet (Thorlabs AC254-200-A-ML). A fluorescein isothiocyanate (FITC) filter set was used for fluorescence imaging (Thorlabs MDF-FITC), and Horiba DeltaDiode DD-485L provided the pulsed illumination for fluorescence excitation. The laser diode emits at 470 nm with a 1-MHz repetition rate; the iCCD was triggered at 500 kHz (on every second pulse).

Samples. The fluorescent bead samples were prepared by drop-casting isopropanol-diluted bead solution on a microscope slide. Once dried, the beads were washed once with isopropanol to wash away weakly adhered beads. Both $2 \mu\text{m}$ (Merck, L4530) and $4 \mu\text{m}$ (Bangs Laboratories Inc., FSDG006) diameter beads were prepared in the same fashion. Acridine orange-stained *Convallaria* samples were obtained from Johannes Lieder GmbH & Co. KG (Catalog number As3212). In literature, Yellow-Green dye ($2 \mu\text{m}$ Merck beads) has been reported

with lifetimes of 2.1 ns (24). Dragon Green dye (4 μm Bangs Laboratory beads) excited at 485 nm has been measured to fluoresce with a 3.4-ns (25) and 4-ns lifetime (26). We validated the lifetime of our samples using a SPAD array (Horiba, FLIMERA); *SI Appendix* for details.

Inverse Retrieval. Our forward model is based on Eqs. 1 and 2. Let us represent as \mathbf{P} the mathematical function that takes as its input the lifetime map τ and the amplitude values A and performs the integration followed by the shifted sum of the replicas as mentioned in Eq. 1; then, we can concisely write the measurement registered on the iCCD sensor as

$$\hat{s} = \mathbf{P}(\tau, A) + n, \quad [5]$$

where n is measurement noise. Integrating Eq. 2 for t from t_0 to ∞ yields $q_{ij} = kA_{ij}\tau_{ij}$. This can be rearranged as $A_{ij} = \frac{q_{ij}}{k\tau_{ij}}$, i.e., the CMOS measurement is used to evaluate the amplitude of decay A_{ij} , which is then fed to the inverse retrieval algorithm. For this, the factor k in Eq. 2 has been calibrated experimentally and is a one-time calibration for the system. The lifetime values, τ , thus become the only unknown variables of function \mathbf{P} henceforth. The cost function we optimize is written in Eq. 4:

$$\hat{\tau} = \arg \min_{\tau} C(\tau), \text{ where} \quad [6]$$

$$C(\tau) = \|\mathbf{P}(\tau) - \hat{s}\|_2 + \alpha \|\tau\|_2^2$$

subject to $\tau \geq 0$.

We initiate the optimization with a random guess for the lifetime map τ . The algorithm then uses gradient descent to reach the optimal solution for the variables τ . If the lifetime map at some iteration n is represented by $\tau^{(n)}$, then the solution at the next iteration progresses as

$$\tau^{(n+1)} = \tau^{(n)} - \beta [\nabla_{\tau} C(\tau)]_{\tau=\tau^{(n)}}. \quad [7]$$

Here, β is the step size in the negative gradient direction and is determined in each iteration by backtracking line search (27). $\nabla_{\tau} C(\tau)$ is the gradient of the cost function with respect to τ :

$$\nabla_{\tau} C(\tau) = \frac{\mathbf{P}(\tau) - \hat{s}}{\|\mathbf{P}(\tau) - \hat{s}\|_2} \frac{\partial \mathbf{P}(\tau)}{\partial \tau} + 2\alpha \sum_{ij} \tau_{ij}. \quad [8]$$

We evaluate the above partial derivative by considering the discrete version of Eq. 1, giving

$$\begin{aligned} \left. \frac{\partial \mathbf{P}(\tau)}{\partial \tau} \right|_{ij} &= -f_1 \sum_t G(t) q_{ij} e^{\frac{t_0-t}{\tau_{ij}}} \left[\frac{t_0-t}{\tau_{ij}^3} + \frac{1}{\tau_{ij}^2} \right] \\ &- f_2 \sum_t G(t) q_{ij-y} e^{\frac{t_0+t_c-t}{\tau_{ij-y}}} \left[\frac{t_0+t_c-t}{\tau_{ij-y}^3} + \frac{1}{\tau_{ij-y}^2} \right] \\ &- \dots \\ &- f_6 \sum_t G(t) q_{ij-5y} e^{\frac{t_0+5t_c-t}{\tau_{ij-5y}}} \left[\frac{t_0+5t_c-t}{\tau_{ij-5y}^3} + \frac{1}{\tau_{ij-5y}^2} \right]. \end{aligned} \quad [9]$$

We use the same notation as in Eqs. 1 and 2.

Convolutional Neural Network. We developed a CNN to retrieve experimental lifetime maps, as an alternative, faster approach to the numerical inversion method described above. The network uses two inputs—the iCCD image and the CMOS image—to predict the lifetime of the sample. We trained the network on a fully synthetic dataset, generated with the forward model in Eq. 5, with model and noise parameters obtained from experimental data. A set of 10,000 270 × 220 pixel intensity and lifetime seed pairs were used to

generate 10,000 220 × 220 single-shot images (the seed images are larger because the single-shot images can have replicas originating from outside their FOV).

The CNN architecture is designed to match the physical parameters of the system and to process the data efficiently. The two inputs are concatenated at the very start. This is because the CMOS data alone cannot provide insight into the lifetime. Therefore, feature extraction layers are superfluous for the CMOS. Similarly, the CMOS constrains the information content of the iCCD.

After concatenation, 256 dilated convolutional kernels are applied, without zero-padding the inputs. Each kernel's size is (11, 1), due to the assumption of 6 replicas, centering the kernel on the 0th replica, with a dilation rate of 10 to match the pixel separation of replicas, creating a sparse receptive field of 101 pixels.

After the first convolutional layer, a second convolutional layer that outputs 64 feature maps is used and then a third that outputs 16 feature maps. The kernel sizes are still (11, 1) with dilation rates of 10. These first three convolutional layers perform what amounts to blind deconvolution, doing much of the processing. Finally, 3 × 3 kernels output the final lifetime map. These 3 × 3 kernels help make the predictions smooth in a local 3 × 3 neighborhood, informing each pixel of the lifetime of its neighbors. This is equivalent to a sliding window binning strategy, except that the window does not use a simple sum, but rather a learned kernel. We find that this reduces image artifacts that might be present if each pixel were treated independently. The output feature map contains the lifetime prediction. We also experimented with predicting lifetime from the iCCD image only, discarding the CMOS image; *SI Appendix* for details.

The network is trained on 7,500 image triplets, validated on 1,250 and tested on 1,250. The training loss we used was mean squared error (MSE) weighted by intensity:

$$L = \frac{1}{N \times M} \sum_i \sum_j q_{ij} (\hat{\tau}_{ij} - \tau_{ij})^2, \quad [10]$$

where N and M are the height and width of the ROI, q is the CMOS field of view corresponding to the lifetime label τ , and $\hat{\tau}$ is the predicted lifetime map. The motivation for scaling MSE is based on the variance of an unbiased lifetime predictor, which is inversely proportional to the total photon count. When measured photon counts are higher, SNR is higher too, and lifetime prediction can be expected to have lower variance; thus, any prediction error has a large significance since it cannot be attributed to a poor SNR and instead indicates poor fitting.

In summary, there are 3 physics-inspired aspects of this neural network. First, the usage of dilated kernels in our convolutional layers mimics the causal relationship between the lifetime prediction at a given pixel and the iCCD pixels above and below this target pixel, as described earlier. Second, the final layer is formed of 3 × 3 kernels, paralleling the sliding window binning strategy commonly used for increasing SNR in noisy measurement. Thirdly, our intensity-weighted mean squared error training loss simulates the variance of an unbiased lifetime predictor, which is inversely proportional to total fluorescent intensity due to the SNR of Poisson distributed photons from the sample (28).

Backpropagation was implemented via adaptive momentum (Adam), using an exponentially decaying learning rate for rapid initial training and eventual fine-tuning. The CNN was trained for 200 epochs, taking ≈ 6 h. Data generation/preprocessing and training/prediction were performed using an Intel(R) i9-9820X CPU and NVIDIA GeForce RTX 2080 Ti, respectively.

Both the preprocessing and CNN are computationally lightweight; for a single input, without special optimization, our code has a wall time of 21.6 ± 0.13 ms and 71.1 ± 6.5 ms, respectively, on the same computer. Hence, the processing pipeline can be run in real-time in parallel with a ≤10-Hz measurement.

Uncertainty Analysis. We analyze the improvement in the theoretical uncertainty of the lifetime prediction for general lifetimes. For this, we use a Bayesian prediction model.

Consider a signal of lifetime τ , creating an integrated CMOS measurement q and a series of noisy iCCD measurements $\hat{s} = (\hat{s}_1, \hat{s}_2, \dots, \hat{s}_3)$. Here, \hat{s}_n is the noisy iCCD signal in the n -th replica. We assume that the CMOS has an SNR far better than the iCCD camera; therefore, we ignore the CMOS in our noise calculations. Hence, we assume that the uncertainties in lifetime

predictions originate from the iCCD only and the probability density function of our predictions likewise.

Continuing Eq. 3, Bayes' theorem states that

$$\begin{aligned} \rho(\tau|\hat{s}) &= \frac{\rho(\hat{s}|\tau)\rho(\tau)}{\rho(\hat{s})} \\ &= \frac{\rho(\hat{s}|\tau)\rho(\tau)}{\int_{\tau} \rho(\hat{s}|\tau)\rho(\tau)d\tau}. \end{aligned} \quad [11]$$

We do not aim to make prior assumptions on $\rho(\tau)$ as we wish the system to be unbiased; however, for the sake of numerical calculation, we limit it between 0.02 and 20 ns. In this equation, $\rho(\hat{s}|\tau)$ is simply dependent on our noise model (for fixed total signal intensity), while the denominator can be numerically estimated by discretizing the integral.

Let us denote the variable denoting replica i 's iCCD signal as S_i , with a value \hat{s}_i in a particular measurement. S_i is distributed as

$$S_i \sim s_i \pm n.$$

We expand Eq. 1 for s_i as

$$s_i = \int_t M(t)D_{QE}P(t)dt, \quad [12]$$

where M is the iCCD gain, D_{QE} is the detector quantum efficiency, and $P(t)$ is the number of photons falling on the pixel at time t and is dependent on τ ; finally, n is the noise (29). Eq. 1 describes the signal in our particular case as a function of lifetime τ .

The iCCD amplifies not only the Poissonian (signal) noise but also dark current noise and clock-induced charge (spurious) noise. Therefore, the noise, n , is

$$\begin{aligned} n &\sim \sqrt{\sigma_{readout}^2 + F^2 \int_t M^2(t)(\sigma_{signal}^2(t) + \sigma_{dark}^2 + \sigma_{cic}^2)dt} \\ \sigma_{signal}^2(t) &= D_{QE} \cdot P(t), \end{aligned} \quad [13]$$

- J. R. Lakowicz, "Time-domain lifetime measurements" in *Principles of Fluorescence Spectroscopy*, J. R. Lakowicz, Ed. (Springer, US, Boston, MA, 2006), pp. 97-155.
- W. Becker, Fluorescence lifetime imaging - techniques and applications. *J. Microsc.* **247**, 119-136 (2012).
- K. Suhling *et al.*, "Fluorescence lifetime imaging" in *Handbook of Photonics for Biomedical Engineering*, A. H.-P. Ho, D. Kim, M. G. Somekh, Eds. (Springer Netherlands, Dordrecht, 2014), pp. 1-50.
- K. Suhling, P. M. W. French, D. Phillips, Time-resolved fluorescence microscopy. *Photochem. Photobiol. Sci.* **4**, 13-22 (2005).
- E. B. van Munster, T. W. J. Gadella, "Fluorescence lifetime imaging microscopy (FLIM)" in *Microscopy Techniques*, J. Rietdorf, Ed. (Springer, Heidelberg, 2005), pp. 143-175.
- H.-J. Lin, P. Herman, J. R. Lakowicz, Fluorescence lifetime-resolved pH imaging of living cells. *Cytometry A* **52A**, 77-89 (2003).
- A. Colom *et al.*, A fluorescent membrane tension probe. *Nat. Chem.* **10**, 1118-1125 (2018). 10.1038/s41557-018-0127-3.
- O. I. Kolenc, K. P. Quinn, Evaluating cell metabolism through autofluorescence imaging of NAD(p)h and FAD. *Antioxid. Redox Sig.* **30**, 875-889 (2017).
- S. Kawanabe, Y. Araki, T. Uchimura, T. Imasaka, Applying fluorescence lifetime imaging microscopy to evaluate the efficacy of anticancer drugs. *Meth. Appl. Fluoresc.* **3**, 025006 (2015).
- Z. Wang *et al.*, Applications of fluorescence lifetime imaging in clinical medicine. *J. Innov. Opt. Health Sci.* **11**, 1830001 (2017).
- J. Jonkman, C. M. Brown, Any way you slice it-a comparison of confocal microscopy techniques. *J. Biomolecular Tech.: JBT* **26**, 54 (2015).
- K. Suhling *et al.*, Wide-field time-correlated single photon counting-based fluorescence lifetime imaging microscopy. *Nucl. Instrum. Methods Phys. Res. Sec. A* **942**, 162365 (2019).
- V. Zickus *et al.*, Fluorescence lifetime imaging with a megapixel SPAD camera and neural network lifetime estimation. *Sci. Rep.* **10**, 1-10 (2020).
- A. Ulku *et al.*, Wide-field time-gated SPAD imager for phasor-based flim applications. *Methods Appl. Fluoresc* **8**, 024002 (2020).
- W. Becker *et al.*, A wide-field TCSPC FLIM system based on a MCP PMT with a delay-line anode. *Rev. Sci. Instrum.* **87**, 093710 (2016).

where F is the noise factor of the amplification process itself. Our Andor iStar 334T uses an 18x-73 intensifier, with quantum efficiency D_{QE} of around 20% at emission above 500 nm, dark current noise of $0.03 e^-/s$ after amplification, readout noise $8 e^-/s$, and clock induced charge noise $0 e^-/s$ (30). We combine Eqs. 12 and 13 as

$$S_i \sim \mathcal{N}(s_i, n). \quad [14]$$

We can then approximate the conditional probability density function $\rho(S_i = \hat{s}_i|\tau)$. We equate our noisy measurement \hat{s} with a noisy lifetime prediction $\hat{\tau}$, giving us $\rho(\hat{\tau}_i|\tau)$. We then multiply the $\rho(\tau)$ (uniform over [0.02,20] ns) and normalize by integrating over τ to obtain $\rho(\tau|\hat{\tau}_i)$ as in Eq. 11

Finally, from $\rho(\tau|\hat{\tau}_i)$, we estimate $\rho(\tau|\hat{\tau})$. For this, we assume that $\hat{\tau}_i$ are measurements of the same underlying variable, which is correct for our given noise model. Therefore,

$$\rho(\tau|\hat{\tau}) = \rho(\tau|\hat{\tau}_0, \hat{\tau}_1, \dots, \hat{\tau}_n) = \prod_i \rho(\tau|\hat{\tau}_i). \quad [15]$$

This probability distribution provides the uncertainty (SD) of estimated lifetime $\hat{\tau}$ for some noisy measurement \hat{s} . This uncertainty is just a lower bound, as our assumptions of a thermal noise-free system reduce the noise of our estimates compared to a physical system. *SI Appendix* for a visualization of our uncertainty analysis.

Data, Materials, and Software Availability. Raw images, synthetic training data (numpy arrays) for training a neural network, code required to process the data and produce the results in the paper, figures in svg and pdf formats. Data have been deposited in Single-shot time-folded fluorescence lifetime imaging (10.5525/gla.researchdata.1380) (31).

ACKNOWLEDGMENTS. We acknowledge funding from EPSRC (UK, grant no. EP/T002123/1). V.K. acknowledges funding from Horiba. V.Z. received funding from European Social Fund (project No 09.3.3-LMT-K-712-23-0132) under a grant agreement with the Research Council of Lithuania. D.F. acknowledges funding from the Royal Academy of Engineering.

- L. M. Hirvonen, K. Suhling, Wide-field TCSPC: Methods and applications. *Meas. Sci. Technol.* **28**, 012003 (2016).
- K. Suhling *et al.*, Wide-field time-correlated single photon counting-based fluorescence lifetime imaging microscopy. *Nucl. Instrum. Methods Phys. Res.* **942**, 162365 (2019).
- R. M. Ballew, J. N. Demas, An error analysis of the rapid lifetime determination method for the evaluation of single exponential decays. *Anal. Chem.* **61**, 30-33 (1989).
- A. J. Bowman, B. B. Klopfer, T. Juffmann, M. A. Kasevich, Electro-optic imaging enables efficient wide-field fluorescence lifetime microscopy. *Nat. Commun.* **10**, 1-8 (2019).
- L. Gao, J. Liang, C. Li, L. V. Wang, Single-shot compressed ultrafast photography at one hundred billion frames per second. *Nature* **516**, 74-77 (2014).
- Y. Ma, Y. Lee, C. Best-Popescu, L. Gao, High-speed compressed-sensing fluorescence lifetime imaging microscopy of live cells. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2004176118 (2021).
- B. Heshmat, M. Tancik, G. Satat, R. Raskar, Photography optics in the time dimension. *Nat. Photo.* **12**, 560-566 (2018).
- K. K. Sharman, A. Periasamy, H. Ashworth, J. N. Demas, Error analysis of the rapid lifetime determination method for double-exponential decays and new windowing schemes. *Anal. Chem.* **71**, 947-952 (1999).
- J. P. Houston, M. A. Naivar, P. Jenkins, J. P. Freyer, Capture of fluorescence decay times by flow cytometry. *Curr. Protoc. Cytom.* **59**, 1-25 (2012).
- M. Patting *et al.*, Fluorescence decay data analysis correcting for detector pulse pile-up at very high count rates. *Opt. Eng.* **57**, 031305 (2018).
- S. Orthaus-Mueller *et al.*, Rapidflim: the new and innovative method for ultra fast flim imaging. *PicoQuant Appl. Note*, 1-8 (2016).
- S. Boyd, S. P. Boyd, L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).
- M. Vitali *et al.*, Wide-field multi-parameter FLIM: Long-term minimal invasive observation of proteins in living cells. *PLoS One* **6**, e15820 (2011).
- Andor. Comparing EMCCD, ICCD and CCD cameras. <https://andor.oxinst.com/learning/view/article/ccd,-emccd-and-iccd-comparisons>, 2023.
- Andor. iStar 334T CCD (2018). <https://andor.oxinst.com/products/intensified-camera-series/istar-334t>.
- V. Kapitanay, A. Fatima, V. Zickus, Single-shot time-folded fluorescence lifetime imaging. Enlighten Research Data. <http://dx.doi.org/10.5525/gla.researchdata.1380>. Deposited 14 March 2023.