

RESEARCH ARTICLE | JULY 22 2025

Convolutional neural network approach to ion Coulomb crystal image analysis

James Allsopp ; Jake Diprose ; Brianna R. Heazlewood ; Chase Zagorec-Marks ;
H. J. Lewandowski ; Lorenzo S. Petralia ; Timothy P. Softley  



J. Chem. Phys. 163, 044201 (2025)

<https://doi.org/10.1063/5.0272967>



View
Online



Export
Citation

Articles You May Be Interested In

A study of the translational temperature dependence of the reaction rate constant between CH_3CN and Ne^+ at low temperatures

J. Chem. Phys. (September 2020)

State selective preparation and nondestructive detection of trapped O^2+

J. Chem. Phys. (February 2025)

Isotope-specific reactions of acetonitrile (CH_3CN) with trapped, translationally cold CCl^+

J. Chem. Phys. (February 2021)



The Journal of Chemical Physics

Special Topics Open for Submissions

[Learn More](#)

Convolutional neural network approach to ion Coulomb crystal image analysis

Cite as: J. Chem. Phys. 163, 044201 (2025); doi: 10.1063/5.0272967

Submitted: 27 March 2025 • Accepted: 30 June 2025 •

Published Online: 22 July 2025



James Allsopp,¹ Jake Diprose,² Brianna R. Heazlewood,² Chase Zagorec-Marks,^{3,4}
H. J. Lewandowski,^{3,4} Lorenzo S. Petralia,⁵ and Timothy P. Softley^{6,a)}

AFFILIATIONS

¹Research Software Group, University of Birmingham, Edgbaston B15 2TT, United Kingdom

²Department of Physics, University of Liverpool, Liverpool L69 7ZE, United Kingdom

³JILA, National Institute of Standards and Technology and the University of Colorado, Boulder, Colorado 80309, USA

⁴Department of Physics, University of Colorado, Boulder, Colorado 80309, USA

⁵Institute for Breath Research, University of Innsbruck, Innsbruck, Austria

⁶School of Chemistry, University of Birmingham, Edgbaston B15 2TT, United Kingdom

^{a)}Author to whom correspondence should be addressed: t.p.softley@bham.ac.uk

ABSTRACT

This paper reports on the use of a convolutional neural network methodology to analyze fluorescence images of calcium-ion Coulomb crystals in the gas phase. A transfer-learning approach is adopted using the publicly available RESNET50 model. It is demonstrated that by retraining the neural network on around 500 000 simulated images, we are able to determine ion-numbers not only for a validation set of 100 000 simulated images but also for experimental calcium-ion images from two different laboratories using a wide range of ion-trap parameters. Absolute ion numbers in the crystal were determined for the experimental data with a percentage error of ~10%. This analysis can be performed in a few seconds for an individual crystal image, and therefore, the method enables the objective, and efficient, analysis of such images in real time. The approach adopted also shows promising performance for identifying Ca⁺ ion numbers in images of mixed-species crystals, thereby enhancing the experimental methodologies for studying the kinetics and dynamics of cold ion–molecule reactions.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution-NonCommercial 4.0 International (CC BY-NC) license (<https://creativecommons.org/licenses/by-nc/4.0/>). <https://doi.org/10.1063/5.0272967>

I. INTRODUCTION

A. Formation and imaging of laser-cooled ionic Coulomb crystals

The use of deep-learning methods to aid analysis of experimental results in the physical sciences has been growing rapidly in recent years. This is particularly true for image analysis,^{1–3} where very large datasets may be generated, and there is a need to extract parameters that characterize the data efficiently and rapidly, even in real time as the data are being recorded. In the current paper, we report the use of a deep-learning method with a convolutional neural network (CNN) to analyze fluorescence images of ion Coulomb crystals (CCs), which are created when ions in the gas phase are laser-cooled in a multipole radio frequency trap (such as a Paul trap⁴). We demonstrate that the efficiency and objectivity of analysis of the images—in terms of deriving the absolute numbers of ions of various

chemical species in a given crystal—can potentially be dramatically enhanced through the use of a CNN. This advance leads to a significant benefit for the experimental methodology of using CCs to study the kinetics and dynamics of cold ion–molecule reactions. We also show that a publicly available CNN model, RESNET50, which was originally trained to categorize images of everyday objects, animals, complex scenes etc., can be retrained to analyze and extract numerical data from scientific images of this type. We anticipate that this approach could be applicable in other imaging applications in chemical physics.

As discussed previously in several articles,^{5–9} ionic CCs are normally formed within a quadrupole (Paul) trap, with the ions created in the trap by ionization of neutral precursor atoms or molecules (e.g., by photoionization or electron-impact ionization), at low densities in an ultrahigh vacuum [Fig. 1(a)]. The trap applies confining




FIG. 1. CC image generation and simulation. (a) Ion trap with imaging set up (calcium oven and effusive beam not shown). (b) (i) Experimental Ca^+ ion fluorescence image of a CC. (ii) Matching simulated image using the molecular dynamics code CCMD (see the text for details). (iii) A simulated 3D representation of the CC, from the CCMD code, corresponding to the image shown in (ii). (c) A sequence of experimental images: (A) pure Ca^+ crystal, (B) “flattened” crystal with Xe^+ sympathetically cooled on the outside, and (C–F) progressive time sequence of images as Xe^+ reacts with NH_3 to form NH_3^+ illustrating the growing hollow core of non-fluorescing NH_3^+ .¹⁰

forces in three dimensions with a combination of radio frequency and electrostatic fields, tending to push the ions together, but the compression of the ion cloud is opposed by the repulsion between like-charged ions. Initially on formation, the packet of ions forms a diffuse, disordered, and mobile cloud because the kinetic energy of the ions exceeds any localized potential energy minima. However, when laser cooling is applied to the ions, most commonly to the alkaline-earth singly charged ions (Ca^+ in the work described here), the temperature can be lowered, typically to the milliKelvin range, where a phase transition occurs to a pseudo-crystalline structure. The ions take up sites on an ordered 3D ellipsoidal multilayered array—a gas phase “crystal”—typically with just a few hundred ions and an ion-to-ion spacing of order $10\text{--}20\ \mu\text{m}$ (cf. the spacing in a true solid-state crystal, which is of order $0.2\ \text{nm}$). The phase change

occurs so as to minimize the potential energy of the system through the ions taking up regular lattice positions at the optimum separation. Figure 1(b-iii) shows a simulated 3D structure of a CC of Ca^+ ions of approximate total length $300\ \mu\text{m}$ on the long axis.

The laser-cooled ions are constantly fluorescing through the optical excitation and de-excitation cycle, and the rate of photon emission and the large (micrometer-scale) separation of the ions are sufficient to enable a simple microscope with an imaging camera to be used to observe layers of the crystal structure with single-ion resolution. Figure 1(b-i) shows an example experimentally recorded Ca^+ ion image, in which each white, diffuse spot represents a single fluorescing Ca^+ ion that is located in the central 2D slice through the CC in the imaging-system focus. Various applications of ionic CCs have been developed in spectroscopy,^{11–13} reaction dynamics, and

kinetics,^{6–9,14–17} and the study of cold ion–molecule collisions has become an important component of the burgeoning field of cold and ultracold chemistry.^{18–20} Such experiments take advantage of the fact that other types of positively charged ions, beyond those that can be laser-cooled, can be sympathetically cooled and trapped in the crystal (e.g., NH_3^+ , C_2H_2^+ , OCS^+ , HN_2^+ , and MgH^+). In general, these other ions are not directly visible in the images because they are not fluorescing under the experimental conditions. However, they tend to either be found as a dark inner shell of the crystal (if they are lighter mass than the laser-cooled Ca^+) or as a dark outer shell (if heavier). Hence, their presence can be detected by the observed deformation of the fluorescing Ca^+ ion framework. Cold ion–molecule reactions can be studied by passing cold, reacting neutral gases through the CC,^{6,7,21–25} or superimposing the ion trap with a trap for neutral atoms or molecules,¹⁷ and observing how the images change with time. Figure 1(c) shows a sequence of images that was recorded over two minutes, in which the initial pure Ca^+ crystal (image A) is “doped” with Xe^+ ions that sit around the outside (note the flattened Ca^+ structure in images B and C). The Xe^+ ions then react with NH_3 (ammonia) by charge transfer to form NH_3^+ , and a “hollow core” appears in the right-hand images (D–F) because the NH_3^+ products of the charge transfer are trapped in the center of the crystal.^{10,26}

The extraction of a reaction rate from such measurements may be achieved for relatively simple reactions (particularly those with only one possible product ion) by recording numerous images of the type shown in Fig. 1(c) as a function of reaction time. The 2D images can then be analyzed to determine the number of ions of each chemical species (mass) in the full 3-dimensional CC structure, at a sequence of time points. One attractive feature of this methodology is that it is a non-destructive detection method, and thus, in principle, a reaction rate can be derived from measurements on a single CC.^{7,24,27,28}

B. Determining ion numbers of CCs

There are a number of reasons why the calcium-ion numbers cannot be routinely read directly from the CC images (unless the crystal contains only a very small number of ions) or determined accurately by simply measuring the overall geometry of the crystal image.

- The images are a two-dimensional cut in the image plane through the 3-dimensional ellipsoidal crystal and do not fully show the detailed 3D structure. Blurred features lying outside the focal plane appear in the crystal images to an extent that depends on the depth of focus in the imaging.
- As the number of ions is increased, the crystal dimensions do not change monotonically with ion number but rather undergo step changes in the dimensions when the number of layers increases. This is particularly the case when there are just a few layers in the crystal. The ion numbers at which these step changes occur vary with the applied fields, trap geometry, and temperature.
- The crystal structure has a dependence on the applied RF field and the endcap voltages, which increasingly squeeze the crystal in the x – y or z dimensions, respectively, as the RF peak-to-peak voltage or endcap voltages are increased. Thus, the number of ions shown in the 2D images changes

as these voltages are adjusted, as well as the size and shape of the overall crystal image.

- In practice, the effective temperature of the ions in the crystal may change experimentally due to a number of possible effects including background collisions with residual gases in the vacuum and imperfect overlap of laser beams. This may lead to blurring of the images and difficulty identifying all the spots in an image.
- The structure is also subject to experimental imperfections due to patch charges on the electrodes or unsymmetrical positioning of the crystal with respect to the trap center, or even due to imperfections on the imaging detector, leading to distortions of the images.

To analyze the images for the number of ions, the crystals can be simulated using an appropriate molecular dynamics (MD) code, which produces simulated images for given ion trap parameters and ion numbers. Figure 1(b-ii) shows the result from a MD simulation of a pure Ca^+ crystal image. Such a simulated image can then be juxtaposed with the experimental image and the number of ions determined by iteration of the initial parameters for the simulation to get a good match. We will refer to this hereafter as the method of *iterative simulated-image matching* (see, for example, Refs. 28–33). Examples of the end result of this matching can be seen later in Fig. 8. In general, the matching of experimental and simulated images is performed by the human eye, rather than an automated procedure. In practice, this is an extremely time-consuming method of analysis—for crystals with a few hundred ions, each MD simulation might take several CPU hours to generate a single simulated image, and the simulation time scales up in proportion to N^2 (where N is the number of ions); then, this comparison of experiment and simulation needs to be iterated. Furthermore, there is inevitably a degree of subjectivity in matching by eye a “perfect” simulated image with an “imperfect” experimental one, especially when trying to define the ion numbers with a percentage error of better than 5%.

An alternative, complementary method for determining ion numbers in the crystal was developed,^{34–37} in which the ions are near-instantaneously ejected from the trap into a time-of-flight mass spectrometer and the number of ions arriving at the detector is measured in mass-selective time gates. While such measurements can be made in real time and indicate directly the mass and relative proportions of all species that are present,^{23,38} one disadvantage of this approach is that the CC is destroyed in the process; therefore, this can be used to determine ion numbers at only a single time point in a reaction kinetics experiment. To make measurements at a number of time points, new crystals must be formed for each time point, and in practice, there will be variations in the initial numbers of ions in each of these crystals that needs to be accounted for to determine the reaction kinetics. In addition, the accuracy of this method relies on calibration of the ion numbers at the detector (and the extraction efficiency), and over a long period of time, the absolute accuracy of ion numbers may change incrementally due to slow degradation of the detector, even if the relative ion-numbers for different masses remain accurate. The collection efficiency of ions of different masses may also differ, especially if a wide spread of ion masses is present,^{38,39} as the ion-extraction field used in a quadrupole ion trap is less homogeneous than in conventional time-of-flight mass spectrometers (see, e.g., Fig. 6 of Ref. 34).

In a recent paper,³³ it was demonstrated that the number of dark ions in a CC could be determined mass-selectively and non-destructively using excitation of the dark-ion secular resonances, a development of the pioneering secular-excitation mass-spectrometry of Ref. 14. Up to 8 H_2^+ ions could be detected quantitatively in a CC of 2000 Be^+ ions, and the chemical change from H_2^+ to H_3^+ was observed resulting from the reaction with background H_2 . However, the step change in fluorescence signal became non-linear for dark-ion counts greater than 8 ions, and the count of Be^+ ions could not be determined by this method (the method of iterative simulated image matching was employed). The method was also limited to dark ions that were of smaller mass than the laser-cooled ions. In some applications, this method could be complementary to the method described in the current paper.

Ideally, one would like to have a real-time, non-destructive method of analysis of an experimental sequence of CC images. In this paper, we report on exploring the use of a deep-learning approach using a CNN to enable rapid and reproducible analysis of ion numbers from CC images. This would bring a greater level of objectivity to the process of extracting ion numbers from an image. It might also allow rapid deduction of experimental rate constants in reaction studies and provide immediate pointers to follow-up measurements that are needed. The method enables determination of absolute ion numbers, which is desirable or essential in many circumstances. While rate coefficients for first-order or pseudo first-order processes can be determined with knowledge of $N_{\text{ion}}(t)/N_{\text{ion}}(t_0)$,⁴⁰ i.e., relative ion numbers, knowledge of absolute ion numbers is needed in situations such as when multiple reaction pathways are present;³⁸ when product ions undergo secondary reactions with reactant gases;³⁴ when one of the products has the same mass as the laser-cooled ions;³⁸ when large CCs are used and detector saturation must be assessed;³⁹ or when background reactions occur with unidentified gases in the vacuum chamber.

We start with the perspective that there is a one-to-one correspondence between the CC image observed and a set of input parameters—the parameters describing the physical geometry and fields applied by the trap and the number of ions of different masses in the crystal. Although that relationship is effectively determined by Newtonian laws of classical physics, we seek to train a neural network to recognize the input parameters—and particularly the number of ions—simply by “looking at” the images, obviating the need to carry out any molecular dynamics simulations for the analysis of experimental images. The inference of ion numbers is immediate—it takes a matter of seconds for each image on a standard laptop computer—and effectively gives ion numbers in “real time.” Furthermore, virtually, all the computational cost is associated with the CNN training process, and once this is completed, the application to infer ion numbers comes at a negligible cost.

Our initial objective was to train a CNN with a sufficiently broad training set (with images generated using a wide range of ion trap parameters), so that it can be used across multiple ion-trap setups and can determine ion numbers from CC images without being provided with any further information about the trapping parameters, or the CCD camera and optics, used to obtain those images. This is analogous to the use of CNNs to identify everyday objects or animals from a photo without being given prior information about the details of the specific camera used, or how close the object was to the camera, or what the resolution of the camera was.

C. Neural networks and their applications

The use of neural networks and artificial intelligence to solve demanding problems has been ubiquitous over the past decade, but the idea of using the human-brain model for a system that can be trained to provide the desired output for a given set of inputs can be traced back to 1943.⁴¹ The first trainable “neural network” with hidden layers was constructed in the 1980s.⁴² Neural networks are a type of supervised learning technique and require an extensive training-dataset, consisting of input data points and their associated outputs. In the current work, the inputs are the CC images in numerical matrix form, and the output data points are parameters such as the ion numbers, or as shown recently by Yin and Willitsch, the secular temperature of the crystal.⁴³

Training takes place by presenting a set of input data points and known associated outputs to the network, which attempts to either classify the input data (for example, identifying whether a photographic image shows a cat or a dog) or to produce values for the output parameters (regression), and this is compared to the known output values, and loss is calculated. The “loss” quantifies the difference between predicted values generated by the CNN model and the actual values. It can be defined as either the cross-entropy in the classification case or the Root Mean Square Error (RMSE) in the regression case. The numerical weights in the network are then adjusted to minimize this loss using the first-order derivative, and the process is repeated. The amount that these weights are altered (with respect to the slope of the loss) is known as the learning rate. In the methodology of the current paper, we adopt the “regression” approach to extract ion numbers, while in the paper of Yin and Willitsch,⁴³ the “classification” approach is adopted.

By 1998, Lecun *et al.*⁴⁴ had developed a system based on the concept of receptive fields, which are used by the human eye to identify features in an image and to make progress in a handwriting analysis problem. These systems became known as CNNs. Layers of CNNs analyze the data on different length scales before the analysis propagates into the traditional dense layers of the network prior to an output classification or output parameter values being produced. A turning point for the practical application of CNNs came when a huge amount of data started to be created by digital cameras, especially in phones, which was easily collated on the Internet. This was coupled with the realization that even consumer Graphical Processing Units (GPUs) could substantially accelerate the training process by orders of magnitude in comparison with using CPUs. The massive amount of training data and the ability to quickly and affordably train networks meant that the architecture of these networks could increase in complexity. CNNs became the leader in this field after the success of AlexNet in the 2010 ImageNet competition, where, across a 1.2×10^6 image sample comprising 1000 classes, it demonstrated an accuracy that was significantly better than the state of the art.⁴⁵ This marked the transition from CNNs being an interesting idea into a practical technique and substantially reduced the computational cost of image recognition.

The objective of the work described in this paper is to build a suitable neural network, or retrain a pre-existing neural network, that can be trained on simulated and/or experimental CC images with known ion numbers and trap/imaging parameters for each image. Having trained the network, we seek to use it to predict the ion numbers for unseen images and without knowledge of other

trapping parameters. Ideally, we aim for this analysis to be achievable in real time while the experiments are taking place—which, in practice, means analysis in a few seconds per image.

II. COMPUTATIONAL AND EXPERIMENTAL METHODS

A. Molecular dynamics simulation and training of the CNN

To solve the problem outlined, we started by constructing a very large dataset of simulated images using the Coulomb Crystal Molecular Dynamics (CCMD) code, described most recently in Ref. 32, on the University of Birmingham's BlueBEAR High-Performance Computing (HPC) facility. The CCMD code simulates the laser cooling of a specified number of ions (typically 100–600 ions in this work), using a set of specified input parameters characterizing the trap dimensions and fields, the imaging parameters, the masses of the ions, the forces due to Coulomb repulsion between ions, laser cooling, and stochastic heating processes (primarily through collisions with background gas). The code is a classical MD routine, rewritten by M. Bell and C. R. Rennick (University of Oxford) in C++, and is based on a pre-existing MD code ProtoMol.⁴⁶ Over a period of more than 20 years, MD simulations have been able to successfully reproduce images of experimental CC images with high accuracy, as illustrated, for example, in Refs. 28–33 and 47 and also in Fig. 8 of the current paper. This is achievable both for pure laser-cooled ion crystals and for mixed-chemical-species crystals.

The force is written as

$$\mathbf{F}_{\text{tot}} = \mathbf{F}_{\text{trap}} + \mathbf{F}_{\text{ion}} + \mathbf{F}_{\text{cool}} + \mathbf{F}_{\text{heat}}, \quad (1)$$

where the terms in the equation represent the trapping forces, the Coulomb repulsion, the laser cooling force, and the stochastic heating force, respectively. The equations of motion are integrated using a velocity-Verlet algorithm across the time cycle of laser cooling and then across a subsequent cycle of image gathering. Typically, each cycle involves around one million time-integration points of 0.2 ns each. This time step is sufficiently small to account for the RF oscillation in the trap (3.8 MHz) and the long-range ion-ion interactions.

The positions and velocities of the particles, sampled over the image-gathering period, are then used to produce a simulated image using the approach illustrated in Fig. 2. The image shows the 2D-sliced probability density of the ions, with each slice corresponding to ions at a specific distance from the focal plane, nominally the center of the crystal. A Gaussian blur is applied to sliced images that are not in the focal plane, with greater blurring for slices that are further from the plane. These slices are then summed together to give a simulated image to compare with experimentally obtained images. The blurring of the outer layers makes them look less intense, but the fluorescence rate and detection efficiency are approximately constant for all ions across the crystal, which is illuminated by a laser with beam waist much larger than the crystal dimensions. We mainly simulate images for pure Ca^+ crystals in this paper, although the CCMD code is capable of simulating images of mixed-chemical-species crystals with a combination of laser-cooled ions and sympathetically cooled ions [such as in Fig. 1(c)], and we made use of that capability as demonstrated later.

For the CNN training set, we chose ranges for the key input parameters that may vary in a range of typical experimental Paul-trap CC studies. The parameters initially varied are as follows:

- 33 ion number values ranging from 30 to 350 in steps of 10.
- 14 radio frequency peak-to-peak voltages ranging from 170 to 300 V in steps of 10 V.
- 21 endcap voltages, ranging from 1.5 to 3.5 V in steps of 0.1 V.
- 24 eta parameters ranging from 0.240 to 0.355 in steps of 0.005 V. This parameter is needed to characterize the imperfection of the field experienced by the ions experimentally; the code adds an additional dc force field, in which eta characterizes the ratio of axial to radial components of the correction field.
- 3 random seeds, used to initiate randomly the velocities and positions of the ions in the simulations. For crystals with several hundred ions, there are several degenerate or near-degenerate structures, some of which are related by symmetry. Using different random seeds helps ensure that these variants are covered in the training set.




FIG. 2. Illustrating the procedure for generation of a simulated 2D image in the CCMD code from the 3D structure produced in the MD simulations. Image created in <https://BioRender.com>.

The heating and cooling rates had constant values for all the simulations. The balance of heating and cooling determines the secular temperature once the crystal reaches equilibrium. Initial testing indicated that crystals did not undergo significant structural changes (in terms of shell structure and numbers of ions in each layer) if the heating:cooling ratio was varied over a reasonable range appropriate for the experiments we were aiming to analyze. Instead, the simulated images showed minor disorder and/or greater blurring as the heating rate increased. A good example of the effect of varying the secular temperature is shown in Fig. 2 of Ref. 43.

Taking all permutations of these values leads to a dataset of 698 544 images. In order to test the quality of the training process, the dataset is partitioned in a 5:1:1 ratio into a training set, a validation set, and a test set, where the “validation” set is used as unseen data to test the accuracy of the predictions of the neural network, and the “test” set is reserved for further testing of the hyperparameters of the model if required. These sets were defined by including each of the parameters in the filename of each image, taking the hash of that filename and taking modulus 7 of that value. The validation and test sets had a modulus of 0 and 1, and this approach had three advantages: it is a simple method; we could start validation before the training simulation had completed without risking leakage between the two datasets; and, most importantly, this should guarantee that there is no bias between the training, validation, and test samples. Ray Tune was used to carry out an unbiased optimization of hyperparameters.⁴⁸

Before training, the simulated images are automatically cropped with a rectangular boundary around the elliptical image in order to reduce the number of black (zero) pixels analyzed by the neural network. We use image augmentation to diversify the data that the network is trained on and to minimize any possibility of over-training, even with the large number of training images. This includes rotating the crystal image, moving it off center, and enlarging or reducing its size. This augmentation is applied randomly before adding a simulated image to the dataset.

Initially, we attempted to build our own CNN, which was coded using PyTorch or PyTorch Lightning. Typically, the models we employed included three CNN layers feeding into three “standard” dense layers performing a regression analysis to produce a single output (the ion number). While these relatively simple models performed moderately well, especially in the ion-number analysis of a validation set of simulated images, they did not perform satisfactorily when given experimental images to analyze.

Subsequently, it was decided to test out the RESNET50 CNN model,⁴⁹ and the remarkable capabilities of that model became quickly apparent. This is the model that was used for all the results presented in this paper. The RESNET50 model was developed by Microsoft Research in 2015 and is a deep CNN with the number “50” referring to the number of layers in the network. The model, which has been applied to diverse applications such as facial recognition and complex-scene analysis, has been trained on the ImageNet dataset that now contains over 15×10^6 images and 1000 classes (images of objects, animals, etc.).⁴⁵ Only a modification of the input and output layers was required in the present work to retrain and use the RESNET50 model for our purposes. For the output layer, we replaced the final classification step, with a conventional dense

numerical output layer, and the predicted ion number was output as a double-precision number. The top layer was changed to work with grayscale images, as well as color images. As shown below, this model has been successfully deployed to analyze data from laboratories in Liverpool/Oxford, UK and Boulder, CO, USA. A self-standing “Inference” routine has been created in this work, which runs on a standard laptop in any location, via a link using the Docker platform to the code. It uses the weights generated in the adapted RESNET50 model, through training on the simulated CC images, to determine the Ca^+ ion number for any CC ion image presented to the routine. The time taken to analyze each image is typically less than 5 s.

B. Generating a test set of experimental CC images

An additional set of experimental CC images was generated to test the accuracy of ion-number prediction by the CNN on real data. These were recorded on the laser-cooled ion-trap time-of-flight mass-spectrometer setup at JILA, University of Colorado Boulder, described in detail in Ref. 35. Ca^+ ions formed via non-resonant ionization of an effusive Ca beam using the third harmonic of a Nd:YAG laser were loaded into a linear Paul trap. The Ca^+ ions were then laser cooled to secular temperatures of <1 K using the frequency-doubled and fundamental outputs of two Ti-sapphire lasers (~ 4 mW at 397 nm and ~ 20 mW at 866 nm). The resultant CC typically consisted of several hundred ions. The fluorescence emitted by the ions of the Ca^+ CC was imaged onto an electron-multiplying charge-coupled device (EMCCD). The number of ions in each crystal was determined via time-of-flight mass spectrometry after the image had been recorded.

A set of 12 000 images was generated, involving creation of 200 individual pure Ca^+ crystals of a range of sizes from ~ 50 to 900 ions. These consisted of 10 sets of 20 crystals with a combination of variable endcap voltages (1.75, 2.5, 3.25, 4.0, and 5 V) and radio frequency peak-to-peak voltages (400 and 250 V). The size (ion number) of each crystal was set approximately by varying the length of time for which the Ca atom beam was ionized. A set of 60 images was recorded for each crystal at 2 s intervals (total observation time: 2 min per crystal), and the ions were then ejected into a time-of-flight mass-spectrometer, and a calibrated Ca^+ ion number was obtained. To confirm this calibration, the method of iterative simulated-image matching was used to compare with the ion numbers derived from the mass spectrum. A sample of some of the images for different crystals obtained at different RF and endcap voltages and ion numbers is shown in Fig. 3(a). A time sequence of images at 2 s intervals is also shown for one crystal in Fig. 3(b).

III. RESULTS

In this section, we present the results of applying the inference routine to determine Ca^+ ion numbers in five different sets of CC images. We begin with the application to the validation set of 100 000 simulated images of pure Ca^+ crystals and then test the level of success the CNN has with simulated mixed-species crystal images (Ca^+ with sympathetically cooled ND_3^+ or CaF^+). It should be noted that the training set included only pure Ca^+ crystals. We then apply inference to *experimental* crystal images, first those recorded




FIG. 3. (a) Experimental images recorded in Boulder for pure Ca^+ crystals with various ion numbers, obtained with an RF peak-to-peak voltage of 400 V (upper line of images) and 250 V (lower images). (b) A sequence of Boulder experimental images recorded for a single CC at 2 s intervals.

in Boulder and then those recorded in Liverpool/Oxford, noting that the CNN was trained on primarily simulated images with the Liverpool/Oxford trap parameters. Finally, we apply it to experimental mixed-species crystals from Liverpool/Oxford with up to 20% of Kr^+ doped into the Ca^+ crystal. These datasets are presented in a progression of what we expected would be an increasing challenge for the CNN.

A. Ion-number inference with simulated pure calcium-ion CC images

The validation-set of $\sim 100\,000$ simulated CC images was run through the CNN validation code to determine the Ca^+ ion numbers, and we compared them with the ion number that was used to generate them. The images were binned into sets according to the actual ion number, with ~ 3000 images in each of 33 bins in step sizes of 10 ions (30–350 ions). Within each bin, the images had widely varying trap parameters as described earlier. The standard deviation of the CNN-inferred ion numbers from the actual ion numbers was calculated for each bin, and these are plotted as the vertical thickness of the pink rectangles in Fig. 4. These deviations are in the range 0.9–3.7 ion-number units or a percentage deviation of 1.2%–1.8%. Example values of the standard deviations in ion numbers are 0.9 at total ion number 50, 2.0 at 150, 2.7 at 250, and 3.7 at 350. Each pink bar is centered vertically on the average value of the ion number produced by the inference routine, and this follows a straight line of slope 1 and passes through the origin. The ten most

outlying points above and below the mean are also shown for each bin in the figure with black circles. A sample of the crystal images represented by the outlying error points is shown in Fig. 5. Although these outlying crystal images are typically almost circular in geometry (spherical crystal) or prolate-shaped (long and thin, pointing vertical), it is not obvious why such categories of image should show




FIG. 4. Plot of inferred Ca^+ ion numbers for 100 000 simulated images (from validation routine) vs the Ca^+ ion number used to simulate each image. The centers of the pink rectangles show the mean inferred ion numbers for the ~ 3000 images in each bin. The vertical thickness (but not the horizontal width) of the rectangle shows the standard deviation for ~ 3000 images. The black open circles show the inferred ion number for the ten greatest outliers in each bin above and below the mean.




FIG. 5. Histogram of output ion numbers from the validation routine for 3040 simulated images with 250 ions. The images for which the deduced ion number has the greatest positive and negative variation from the actual ion number are shown above and below the histogram.

higher errors, except perhaps that the RF and endcap voltages used to generate these lie at the extremes of the range incorporated in the training set—and hence, the CNN has not fully recognized the trends in images derived from that range of parameters. In summary, the CNN is able to predict ion numbers of unseen simulated CC images, generated with a wide spread of trap parameters, with a percentage error of 1%–2% in a few seconds per image.

B. Analysis of simulated mixed-species CC images (Ca^+/CaF^+ and $\text{Ca}^+/\text{ND}_3^+$)

An intriguing question is whether the neural network can determine calcium-ion numbers in images of mixed-chemical-identity crystals. As explained earlier, ions of higher mass than calcium—e.g., calcium fluoride CaF^+ , mass 59 u, tend to take positions in layers around the edge of the crystal, and the fluorescing calcium-ion framework takes on the appearance of a flattened ellipsoid. With this very different shape from the pure calcium crystal, it might be expected that a CNN that has not been trained to recognize such crystals would have difficulty in generating physically reasonable answers for the calcium-ion numbers.

A test set of simulated mixed-crystal images was generated with combinations of 150–250 Ca^+ ions and 0 to 100 CaF^+ ions; a sample of these is shown in Fig. 6(a). The application of the inference routine to deduce the Ca^+ numbers for this test set of 36 images is shown in Fig. 6(b). It is clear that even for this small number of images tested, there is a good prediction of Ca^+ ion numbers with an RMS error of 6.4 and an RMS percentage error of 3.1%.

In the case of mixed crystals with low-mass sympathetically cooled ions, the images give the appearance of having a hollow core. This is illustrated for a set of simulated images for Ca^+ with ND_3^+ (mass 20 u) shown in Fig. 7(a). Given that the neural network had not been trained on such hollow-core crystals, we expected that the prediction of calcium-ion numbers from such images might be more difficult than for the flattened images in Fig. 6. Figure 7(b) shows the outcome of predicted ion numbers for 36 crystals of $\text{Ca}^+/\text{ND}_3^+$ composition with combinations of 150–250 Ca^+ ions and 0 to 100 ND_3^+ ions. The RMS ion number error is 22 for this set, and the RMS percentage error is 11%. While the accuracy of prediction is not as good as for the set in Fig. 6, we consider the accuracy to be reasonable, as is the ability to get the correct ordering of the Ca^+




FIG. 6. (a) Simulated crystals of Ca^+ , which contain a proportion of sympathetically cooled CaF^+ . Each crystal has 250 ions in total and CaF^+ ion numbers ranging from 0 to 100 as indicated. (b) Inference results, showing deduced ion number vs actual Ca^+ ion number for simulated Ca^+/CaF^+ crystals with combinations of 0–100 CaF^+ ions and 150–250 Ca^+ ions.




FIG. 7. (a) Simulated crystals of Ca^+ , which contain a proportion of sympathetically cooled ND_3^+ . Each crystal has 250 ions in total and ND_3^+ ion numbers ranging from 0 to 100 as indicated. (b) Inference results, showing deduced ion number vs actual Ca^+ ion number for simulated $\text{Ca}^+/\text{ND}_3^+$ crystals with 0–100 ND_3^+ ions and 150–250 Ca^+ ions.

ion numbers in general. Furthermore, the slope of the trend is still ~ 1 , and only the intercept is shifted. The tendency for the model to over-estimate the ion numbers is perhaps not surprising given that the size of the ellipse encompassing the ions is somewhat larger than the CNN would “expect” for a crystal with that number of Ca^+ ions. In principle, if this overestimation could be demonstrated to be systematic, then it could be corrected, enabling use of the CNN for identifying numbers in such mixed crystals with light sympathetically cooled ions. As discussed later, the CNN could also be retrained to include a set of simulated mixed-crystal images in the training set.

C. Analysis of experimental pure Ca^+ CC images recorded in Boulder

A test set of 370 experimental images, recorded as described in Sec. II B, was selected from the 12 000 images available. We chose to focus on the images recorded with 250 V RF peak-to-peak voltage because the 400 V RF peak-to-peak voltage value was beyond the range of parameters used in our training set, and it also generated a number of long, thin, vertically aligned crystals that pushed toward, or over, the boundaries of the imaging system. For these 400 V-RF

images, the inference routine tends to significantly underestimate the ion numbers at the higher end of the ion number range (more than 350 ions). It is likely that this would be corrected by the use of different focal length lens experimentally to compress the image size and ensure that it is comfortably within the bounds of the imaging detector.

From each of the five sets of 20 crystals for the 250 V RF voltage (with endcap voltages 1.75, 2.5, 3.25, 4.0, and 5.0 V), we selected 6–10 crystals from each set that covered the range from the smallest crystal size to an ion number of ~ 450 at rough intervals of 50. It should be noted that it is not possible experimentally to prepare a crystal with a precise number of ions on demand. We used the first ten images (out of 60) recorded at two-second intervals for each of the 37 separate crystals.

Prior to running through the inference routine, the experimental images are converted so that they have the same magnification and pixel array dimensions as used in the training set (which was based on the pixel array dimensions for Liverpool images). The adjustment of the magnification factor also takes into account the slightly different trap geometry; for the Liverpool/Oxford trap, $r_0 = 3.5$ mm, while for the Boulder trap, $r_0 = 3.91$ mm, where r_0 is

the minimum distance of each cylindrical electrode from the trap center. Crystals were rotated by $\sim 5^\circ$ to bring the ellipse axis into the vertical orientation, and any non-zero pixels outside the elliptical boundary of the fluorescence from the crystal are set to zero. A five-point two-dimensional median averaging was also applied to reduce the noise level.

Figures 8(a) and 8(b) show examples of the plot of CNN-predicted ion number, for crystals obtained using a common endcap voltage (3.25 and 4.0 V respectively), against the ion numbers determined by the method of iterative simulated-image matching (an example of the outcome of such matching is shown in the inset of each figure). We have chosen to use the iterative simulated-image matching ion numbers for the horizontal axis in these plots, rather than the mass-spectrometer numbers, to enable a fair comparison with results for the Liverpool/Oxford data (where we do not have calibrated mass-spectrometric ion numbers currently.) Table I lists the RMS- and mean-percentage errors for all five datasets. It can be seen that the RMS error is typically around 10%. A significant part of that error stems from the dispersion of values determined for the ten images of each crystal—as is apparent from Fig. 8. If the mean ion-number prediction is used for those ten images in place of the

TABLE I. Mean and RMS errors for inference predictions of Ca^+ ion numbers in five sets of images recorded in Boulder.

End cap voltage/V	Number of images	Ion number range	Mean % error	RMS % error
1.75	8×10	110–411	11.6	15.3
2.5	6×10	115–450	9.2	12.3
3.25	10×10	105–470	7.3	9.9
4.0	7×10	120–440	7.5	10.2
5.0	6×10	115–392	7.2	8.9

individual values for the 3.5 V set, the mean and RMS percentage errors reduce to 5.3% and 6.4%, respectively. If it were possible to identify in advance which of the ten images in each of these sets would give the best predicted ion numbers, then the errors could be reduced to around 1%. However, inspection of the crystal images does not give any obvious grounds for making that selection. Nevertheless, given that the CNN has not used any experimental images in its training, and when considering all the imperfections of such




FIG. 8. Results for Ca^+ crystals recorded in Boulder, showing deduced ion number from inference vs Ca^+ ion number derived by the method of iterative simulated-image matching [see the insets for comparison of example simulated image (green) with observed image (BW)]. (a) The results for ten images recorded at 2 s intervals are shown for each of ten crystals. The ion trap parameters are as follows: RF peak-to-peak voltage 250 V; endcap voltage 3.25 V. A line of slope 1 passing through the origin is shown. (b) Same as in (a) but with seven crystals recorded at an endcap voltage of 4.0 V.

experimental data, we believe that this is an encouraging result. It must also be recognized that there is an “error,” of up to 5%, in the ion numbers determined by the iterative simulated-image matching method, and these are assumed here to be the “correct” numbers for calculating the errors. In practice, an exact match of simulations and experimental images is rarely achievable, given that there are inevitably many near-degenerate structures for the crystal at the crystal sizes and temperatures used here (and the CCMD simulation will land on one of these) and also the imperfections of experimental images play a role in contributing to the error.

D. Analysis of pure Ca^+ crystal images recorded in Liverpool/Oxford

To test the transferability of our trained model across experimental setups and trap parameters, the inference routine has also been applied to deduce ion numbers for a set of 90 images that were recorded (see Refs. 10 and 26 for further details) on an ion-trap setup in the Heazlewood Group in Oxford or Liverpool during the period 2019–2022, on 18 separate days. Compared to the Boulder experiments, the physical dimensions of the Liverpool/Oxford ion trap are slightly different, the cooling laser power is lower,




FIG. 9. (a) Example experimental images of Ca^+ crystals from Liverpool/Oxford, recorded in the period 2019–2022 (unpublished). (b) Example of the effect of the pre-manipulation of experimental images from Liverpool/Oxford as described in the main text. (c) Inference results for pre-manipulated Ca^+ crystal experimental images recorded in Liverpool/Oxford, showing deduced ion number from inference vs Ca^+ ion number derived by iterative simulated-image matching. A line of slope 1 passing through the origin is shown.

and a different camera is used with lower resolution. For each of the images, the ion number had been previously deduced by the method of iterative simulated-image matching. The trap parameters used to record these images were a radio frequency peak-to-peak voltage $V_{RF} = 200\text{--}210$ V and endcap voltage $V_{end} = 2.25\text{--}2.95$ V, and deduced ion numbers ranged from 150 to 720 ions. As illustrated by the sample of images in Fig. 9(a), the sharpness of the images and signal-to-noise level are more variable than the images recorded in Boulder, principally as a consequence of the older imaging camera deployed, and there is a tendency for the outer elliptical layer to appear less intense than we would expect from simulations. In order to increase the similarity of the images to the simulated images on which the CNN had been trained, we performed some additional pre-processing steps compared to the Boulder-recorded images. In particular, the intensity around the outer layer was artificially enhanced to improve uniformity of intensity using a square function of the ellipse equation $[f(x, y)]$ which is equal to 1 at the ellipse boundary],

$$I(x, y) = I_0(x, y) + A \times I_0(x, y) \times (f(x, y))^2, \quad (2)$$

where

$$f(x, y) = \frac{(x-h)^2}{a^2} + \frac{(y-k)^2}{b^2}, \quad (3)$$

and the ellipse is centered at co-ordinates (h, k) and has major and minor axes of length $2a$ and $2b$. $I_0(x, y)$ is the unmanipulated intensity, and A is a constant that had an identical magnitude for all images. In addition, the intensity of images is normalized, such that the “brightness density” = sum of non-zero pixels/number of non-zero pixels is constant across the image set. The overall effect of these subtle changes is illustrated for one crystal in Fig. 9(b).

Figure 9(c) shows a plot of the deduced Ca^+ ion number from the inference routine vs that deduced previously by the method of iterative simulated-image matching. Given that the images are significantly more variable in quality than for the Boulder set, this level of accuracy and the matching of the general trend of the data to a straight line of slope 1 passing through the origin are, in our view, satisfying. The mean percentage error in ion-number prediction for this dataset is 10.4%, which is comparable with that achieved for the Boulder data.

E. Analysis of mixed Ca^+/Kr^+ crystal images recorded in Liverpool/Oxford

Finally, we also tried applying the inference routine to a set of experimental mixed-crystal images, analogous to those shown in Fig. 6. Figure 10(a) shows four examples of the crystal images that were used. These were obtained in a series of experiments conducted in Oxford in 2020.^{26,50} The selected images were all mixed crystals of Ca^+/Kr^+ . We eliminated images from the sample that contained greater than 500 calcium ions, those that contained more than 20% rare-gas ions, and those that appeared visually to be highly asymmetric in the vertical direction.

Figure 10(b) shows a plot of the Ca^+ ion number deduced by the CNN inference routine for mixed crystals of Ca^+/Kr^+ vs the ion number deduced by iterative simulated-image matching. The RMS-percentage deviation of Ca^+ ion number is 12.5%, and the mean percentage deviation is 9.3%. We are not able to deduce the rare-gas ion numbers using the CNN because it has not been trained




FIG. 10. (a) Example experimental images of mixed-species Ca^+/Kr^+ crystals from Liverpool/Oxford, recorded as described in Refs. 26 and 50 in the period 2019–2022. (b) Results for pre-manipulated Ca^+/Kr^+ crystal experimental images recorded in Liverpool/Oxford, showing deduced ion number from inference vs Ca^+ ion number derived by iterative simulated-image matching. The crystals have varying numbers of Kr^+ ions that are less than 20% of the total ion number. A line of slope 1 passing through the origin is shown.

to do that. In principle, it would be possible to retrain the CNN with mixed crystals and to train it to deduce both the Ca^+ ion number and the dopant (e.g., rare-gas ion) number. However, we have not performed that training to date.

Nevertheless, this limited exercise demonstrates once again the versatility of the model that can be extended to make ion-number predictions for experimentally imperfect mixed crystal images (Ca^+/Kr^+) of the type it had not seen in training.

IV. CONCLUSION AND FUTURE PERSPECTIVE

Our overall objective at the outset of this project was to develop (or retrain) a global CNN model that was capable of determining ion numbers from CC images in real time during experimental recordings, in order to enhance the methodology of using CCs for cold ion–molecule reaction dynamics and kinetics. We aimed to train the neural network on a sufficiently diverse dataset, enabling it to be used for images obtained in different laboratories without retraining. This would ensure adaptability to variations in quadrupole ion trap geometries, applied fields, and, ultimately, mixed-species ion crystals. The work reported in this paper demonstrates that this goal is a challenging, but potentially achievable, objective. We have successfully used a slightly modified CNN model (RESNET50) that has been retrained on a very large set of simulated images, which were

generated primarily using the trap parameters of the Liverpool/Oxford laboratory, and applied it with some success to experimental results from both that laboratory and from Boulder. This was achieved without inclusion of experimental data in the training set. The typical errors of $\sim 10\%$ are larger than estimated standard deviations from mass-spectrometry or the method of iterative simulated-image matching ($\sim 5\%$). However, in the case of the Boulder images, where we had multiple images of the same crystal available, it was shown that using the average predicted ion number for ten images led to an improvement in accuracy with a standard deviation of around 6% . Moreover, there are some key potential advantages in the CNN approach, including the non-destructive nature of the ion-number determination (in contrast to the mass-spectrometric approach) and the fast analysis of images—in seconds, not days—in contrast to the iterative simulated-image matching. Furthermore, we have shown that, even though we trained the CNN only on simulated images of single-component crystals (pure Ca^+ crystals), it could be applied in limited circumstances to identify the Ca^+ ion numbers for two-component crystals, especially those where the second component was a species heavier than Ca^+ , for both simulated and experimental images.

Although this is a very promising outcome, the challenges of generating and working with a training set of more than half a million images should not be underestimated. The training set took many months to generate on the University of Birmingham's high-performance computing setup (BlueBEAR). Ideally, we would follow up the present work by generating an extended training set on which to retrain the model, with inclusion of simulated images of multi-component crystals to add to the existing training set, to improve the accuracy and versatility of identifying ion numbers in such crystals. We would also generate enough experimental images, with pre-determined ion numbers (via another method), to add these to the training set. However, we would probably need to have at least 10% of the total images in the training set being experimentally generated to have any impact on the training (i.e., $50\,000$ – $100\,000$ such images). Such large sets of ion-number-labeled experimental images do not exist currently.

However, even with current levels of accuracy, the approach described here could be used immediately in a hybrid methodology for identifying the ion numbers, in parallel with other measurements. In the case of mass-spectrometric analysis, the CNN-based evaluation of ion numbers would be extremely useful for routine checking of the calibration of calcium-ion peak heights as a function of crystal size. In the case of the iterative simulated-image matching, the CNN inference ion number could provide an excellent starting point for further refinement through image matching even for mixed-species crystals. We also note that a typical sequence in chemical kinetic measurements, as shown in Fig. 1(c), always starts with a pure calcium crystal, and knowing the initial Ca^+ number can streamline the process of image simulation of the mixed crystals. Furthermore, the very large database of $700\,000$ simulated Ca^+ crystals forms a valuable, diverse library of images, which could be used to circumvent the need to simulate new images for image matching within the parameter ranges covered in the database.

In other complementary work conducted in Basel in parallel with this work, a more targeted CNN approach was adopted.⁴³ A smaller training set (7200 images) was created, consisting of simulated CC images generated with a relatively narrow set of ion-trap

parameters and using the RESNET 18 and ALEXnet models. Ca^+ ion numbers and secular crystal temperatures were deduced from simulated and experimental images of pure Ca^+ crystals. When applied to a validation set of simulated Ca^+ images with ion number 100 – 299 and fixed ion trap parameters, the RESNET 18 model inferred the correct ion number for 93% of the tested images (classified in steps of one ion-number unit). In some circumstances, this targeted approach may be advantageous where a fixed set of ion trap parameters is used extensively and repeatably in a particular experimental ion trap. However, the use of a targeted training regime implies that if the ion-trap parameters are varied significantly outside that narrow range (as they frequently are in the Boulder setup for example), then a completely new training set would most likely need to be generated each time, and the model may not easily be transferred across different experimental setups without retraining and validation.

Finally, we observe that the model that was eventually retrained and used in the current work—the RESNET50 model—was not originally developed for image recognition of scientific images, but rather for recognition/characterization of images of everyday objects, animals, facial recognition, and complex-scene analysis. The reason why it works also for these scientific images is that, in both cases, the different layers of the CNN are used to characterize features of the image at different length scales, before compiling the information to achieve an overall recognition. This multi-lengthscale characterization is needed in the process to recognize a CC image with a given ion number, as much as it is to recognize an image of a human face. It seems likely that this could be a versatile tool for a wide range of other imaging applications in the physical sciences, as has already been demonstrated in the area of medical imaging.⁵¹

ACKNOWLEDGMENTS

T.P.S. is grateful to the Leverhulme Trust for Grant No. EM-2022-027 to support this work, for a Visiting Fellowship at JILA, University of Colorado Boulder, in 2023, and for the expertise and support of the University of Birmingham Research Software Engineering Group. B.R.H. is grateful to the Engineering and Physical Sciences Research Council (EPSRC Project No. EP/N032950/2), the European Commission (ERC Starting Grant Project No. 948373), and the Leverhulme Trust (Project Nos. RPG-2022-265 and PLP-2022-215) for funding. H.J.L. and C.Z.-M. note support from the US National Science Foundation (Grant No. PHY-2317149) and the US Air Force Office of Science and Research (Grant No. FA9550-20-1-0323). We are grateful to Yanning Yin and Stefan Willitsch (Basel) for sharing their draft manuscript (Ref. 43) in advance of publication.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

James Allsopp: Conceptualization (equal); Data curation (lead); Formal analysis (lead); Funding acquisition (lead); Investigation (lead); Methodology (lead); Project administration (lead); Resources

(equal); Software (lead); Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (supporting); Writing – review & editing (equal). **Jake Diprose:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – review & editing (equal). **Brianna R. Heazlewood:** Conceptualization (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – review & editing (equal). **Chase Zagorec-Marks:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Validation (equal); Visualization (equal); Writing – review & editing (equal). **H. J. Lewandowski:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Funding acquisition (equal); Investigation (equal); Methodology (equal); Project administration (equal); Resources (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – review & editing (equal). **Lorenzo S. Petralia:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Validation (equal); Visualization (equal); Writing – review & editing (equal). **Timothy P. Softley:** Conceptualization (lead); Data curation (equal); Formal analysis (equal); Funding acquisition (lead); Investigation (equal); Methodology (equal); Project administration (lead); Resources (equal); Software (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (lead); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- ¹Z. Zhu, J. Lu, S. Yuan, Y. He, F. Zheng, H. Jiang, Y. Yan, and Q. Sun, *J. Phys. Chem. Lett.* **15**, 1985 (2024).
- ²M. Ziatdinov, A. Ghosh, C. Y. Wong, and S. V. Kalinin, *Nat. Mach. Intell.* **4**, 1101 (2022).
- ³I. Sokolov, *Phys. Chem. Chem. Phys.* **26**, 11263 (2024).
- ⁴W. Paul, *Rev. Mod. Phys.* **62**, 531 (1990).
- ⁵M. Drewsen, I. Jensen, J. Lindballe, N. Nissen, R. Martinussen, A. Mortensen, P. Staunum, and D. Voigt, *Int. J. Mass Spectrom.* **229**, 83 (2003).
- ⁶S. Willitsch, *Int. Rev. Phys. Chem.* **31**, 175 (2012).
- ⁷B. R. Heazlewood and T. P. Softley, *Annu. Rev. Phys. Chem.* **66**, 475 (2015).
- ⁸D. Zhang and S. Willitsch, *Cold Chemistry: Molecular Scattering and Reactivity Near Absolute Zero* (RSC Publishing, 2017), Chap. 10, p. 496.
- ⁹B. R. Heazlewood and H. J. Lewandowski, *Emerging Trends in Chemical Applications of Lasers, ACS Symposium Series* (American Chemical Society, 2021), Vol. 1398, Chap. 17, p. 389.
- ¹⁰L. S. Petralia, A. Tsikritea, J. Loreau, T. P. Softley, and B. R. Heazlewood, *Nat. Commun.* **11**, 173 (2020).
- ¹¹I. V. Kortunov, S. Alighanbari, M. G. Hansen, G. S. Giri, V. I. Korobov, and S. Schiller, *Nat. Phys.* **17**, 569 (2021).
- ¹²M. Germann, X. Tong, and S. Willitsch, *Nat. Phys.* **10**, 820 (2014).
- ¹³P. Staunum, K. Højbjerg, P. Skyt, A. Hansen, and M. Drewsen, *Nat. Phys.* **6**, 271 (2010).
- ¹⁴T. Baba and I. Waki, *J. Chem. Phys.* **116**, 1858 (2002).
- ¹⁵K. Molhave and M. Drewsen, *Phys. Rev. A* **62**, 011401 (2000).
- ¹⁶B. Roth, P. Blythe, H. Wenz, H. Daerr, and S. Schiller, *Phys. Rev. A* **73**, 042712 (2006).
- ¹⁷M. Deiß, S. Willitsch, and J. H. Denschlag, *Nat. Phys.* **20**, 713 (2024).
- ¹⁸J. Toscano, H. J. Lewandowski, and B. R. Heazlewood, *Phys. Chem. Chem. Phys.* **22**, 9180 (2020).
- ¹⁹T. Softley, *Proc. R. Soc. London, Ser. A* **479**, 20220806 (2024).
- ²⁰B. R. Heazlewood and T. P. Softley, *Nat. Rev. Chem.* **5**, 125 (2021).
- ²¹B. R. Heazlewood, *Mol. Phys.* **117**, 1934 (2019).
- ²²A. Kilaj, J. Wang, P. Straňák, M. Schwilk, U. Rivero, L. Xu, O. A. von Lilienfeld, J. Küpper, and S. Willitsch, *Nat. Commun.* **12**, 6047 (2021).
- ²³G. S. Kocheril, C. Zagorec-Marks, and H. J. Lewandowski, *Nat. Astron.* **9**, 685 (2025).
- ²⁴S. Willitsch, M. T. Bell, A. D. Gingell, S. R. Procter, and T. P. Softley, *Phys. Rev. Lett.* **100**, 043203 (2008).
- ²⁵T. Yang, A. Li, G. K. Chen, Q. Yao, A. G. Suits, H. Guo, E. R. Hudson, and W. C. Campbell, *Sci. Adv.* **7**, eabe4080 (2021).
- ²⁶A. Tsikritea, K. Park, P. Bertier, J. Loreau, T. P. Softley, and B. R. Heazlewood, *Chem. Sci.* **12**, 10005 (2021).
- ²⁷A. D. Gingell, M. T. Bell, J. M. Oldham, T. P. Softley, and J. N. Harvey, *J. Chem. Phys.* **133**, 194302 (2010).
- ²⁸M. T. Bell, A. D. Gingell, J. M. Oldham, T. P. Softley, and S. Willitsch, *Faraday Discuss.* **142**, 73 (2009).
- ²⁹P. Blythe, B. Roth, U. Frohlich, H. Wenz, and S. Schiller, *Phys. Rev. Lett.* **95**, 183002 (2005).
- ³⁰B. Roth, A. Ostendorf, H. Wenz, and S. Schiller, *J. Phys. B: At. Mol. Opt. Phys.* **38**, 3673 (2005).
- ³¹C. B. Zhang, D. Offenberger, B. Roth, M. A. Wilson, and S. Schiller, *Phys. Rev. A* **76**, 012719 (2007).
- ³²J. Oldham, D.Phil. thesis, University of Oxford, 2014.
- ³³F. Schmid, J. Weitenberg, J. Moreno, T. W. Hänsch, T. Udem, and A. Ozawa, *Phys. Rev. A* **106**, L041101 (2022).
- ³⁴K. A. E. Meyer, L. L. Pollum, L. S. Petralia, A. Tauschinsky, C. J. Rennick, T. P. Softley, and B. R. Heazlewood, *J. Phys. Chem. A* **119**, 12449 (2015).
- ³⁵P. C. Schmid, J. Greenberg, M. I. Miller, K. Loeffler, and H. J. Lewandowski, *Rev. Sci. Instrum.* **88**, 123107 (2017).
- ³⁶C. Schneider, S. J. Schowalter, K. Chen, S. T. Sullivan, and E. R. Hudson, *Phys. Rev. Appl.* **2**, 034013 (2014).
- ³⁷D. Rösch, H. Gao, A. Kilaj, and S. Willitsch, *EPJ Tech. Instrum.* **3**, 5 (2016).
- ³⁸C. Zagorec-Marks, G. S. Kocheril, O. A. Krohn, T. Kieft, A. Karpinska, T. P. Softley, and H. J. Lewandowski, *Faraday Discuss.* **251**, 125 (2024).
- ³⁹J. A. Diprose, V. Richardson, P. Regan, A. Roberts, S. Burdin, A. Tsikritea, K. Mavrokoridis, and B. R. Heazlewood, *J. Phys. Chem. A* **128**, 3900 (2024).
- ⁴⁰S. Willitsch, M. T. Bell, A. D. Gingell, and T. P. Softley, *Phys. Chem. Chem. Phys.* **10**, 7200 (2008).
- ⁴¹W. S. McCulloch and W. Pitts, *Bull. Math. Biophys.* **5**, 115 (1943).
- ⁴²D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Nature* **323**, 533 (1986).
- ⁴³Y. Yin and S. Willitsch, *arXiv:2502.18442* (2025).
- ⁴⁴Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, *Proc. IEEE* **86**, 2278 (1998).
- ⁴⁵A. Krizhevsky, I. Sutskever, and G. E. Hinton, *Commun. ACM* **60**, 84 (2017).
- ⁴⁶T. Matthey, T. Cickovski, S. Hampton, A. Ko, Q. Ma, M. Nyerges, T. Raeder, T. Slabach, and J. A. Izaguirre, *ACM Trans. Math. Software* **30**, 237 (2004).
- ⁴⁷F. H. J. Hall and S. Willitsch, *Phys. Rev. Lett.* **109**, 233202 (2012).
- ⁴⁸R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica, *arXiv:1807.05118* (2018).
- ⁴⁹<https://en.innovatiana.com/post/discover-resnet-50> for a description of RESNET-50 and for guidance on how to adapt it to new data sets via Transfer Learning, 2024.
- ⁵⁰A. Tsikritea, J. A. Diprose, J. Loreau, and B. R. Heazlewood, *ACS Phys. Chem. Au* **2**, 199 (2022).
- ⁵¹W. Xu, Y.-L. Fu, and D. Zhu, *Comput. Methods Programs Biomed.* **240**, 107660 (2023).