

PAPER • OPEN ACCESS

Machine-learning enhanced dark soliton detection in Bose–Einstein condensates

To cite this article: Shangjie Guo *et al* 2021 *Mach. Learn.: Sci. Technol.* **2** 035020

View the [article online](#) for updates and enhancements.



PAPER

OPEN ACCESS



RECEIVED
31 December 2020REVISED
19 February 2021ACCEPTED FOR PUBLICATION
9 March 2021PUBLISHED
15 June 2021

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Machine-learning enhanced dark soliton detection in Bose–Einstein condensates

Shangjie Guo¹, Amilson R Fritsch¹, Craig Greenberg², I B Spielman¹  and Justyna P Zwolak^{2,*} ¹ Joint Quantum Institute, National Institute of Standards and Technology, and University of Maryland, Gaithersburg, MD 20899, United States of America² National Institute of Standards and Technology, Gaithersburg, MD 20899, United States of America

* Author to whom any correspondence should be addressed.

E-mail: jpzwolak@nist.gov**Keywords:** dark solitons, machine learning, convolutional neural network

Abstract


Most data in cold-atom experiments comes from images, the analysis of which is limited by our preconceptions of the patterns that could be present in the data. We focus on the well-defined case of detecting dark solitons—appearing as local density depletions in a Bose–Einstein condensate (BEC)—using a methodology that is extensible to the general task of pattern recognition in images of cold atoms. Studying soliton dynamics over a wide range of parameters requires the analysis of large datasets, making the existing human-inspection-based methodology a significant bottleneck. Here we describe an automated classification and positioning system for identifying localized excitations in atomic BECs utilizing deep convolutional neural networks to eliminate the need for human image examination. Furthermore, we openly publish our labeled dataset of dark solitons, the first of its kind, for further machine learning research.

1. Introduction

Machine-learning (ML)-based image classification has found application throughout science, from analysis of experimental data in particle physics [1–3], dark matter search experiments [4, 5] or quantum dots experiments [6–9] to predicting properties of materials [10–12] to studying molecular representations and properties [13–15]. In atomic physics, ML has been used to locate topological phase transitions [16], to complement absorption imaging technique [17], to characterize particles in disordered fields [18], and to detect quantum vortices in rotating BECs [19]. In this paper, by combining convolutional neural networks (ConvNets) with traditional fitting techniques, we first categorize many-body atomic physics data, and then extract quantitative information from this data.

Using cold-atom Bose–Einstein condensates (BECs), we focus on solitons, robust solitary waves that retain their size, shape, and speed at which they travel [20, 21]. These properties arise from an interplay between nonlinearity and dispersion that is present in many physical systems. Indeed, since their first observation in canals [22], solitons have been found in rivers and seas [23, 24]; BECs [25, 26]; optical fibers [27, 28]; astronomical plasmas [29]; and even human blood vesicles [30, 31]. Due to their inherent stability, solitons in optical fibers [32] have found commercial applications in long-distance, high-speed transmission lines [33].

While the natural environment does not allow for the controlled study of quantum solitons, BECs are an excellent medium where individual or multiple solitons can be created on-demand, with all their properties, such as position and velocity, tuned according to necessity [34, 35]. Most measurements in BEC experiments produce raw data in the form of images that, in our context, provide information about the solitons' positions within the BEC. The challenge is to efficiently and reliably identify the number of solitons and their locations. Traditional least-squares fitting techniques can locate solitons, provided that the soliton number is known in advance. Currently, the number of solitons is determined manually [35], and this human intervention inhibits the automated analysis of large datasets.



Here, we describe our reliable automated soliton detection and positioning system that takes as input image data and outputs information whether a single soliton is present, and, if so, its location. Since solitons are easily identifiable by human examination of images, this problem naturally connects to the field of computer vision and ConvNet-based image classification [36]. Our algorithm consists of a data preprocessor that converts raw data into a ConvNet-compatible format; a ConvNet image classifier that determines if a single soliton has been detected; and a position regressor that locates the soliton within the BEC, when applicable (see figure 1 for a schematic of the analysis flow).

We show that our fully automated system performs comparably to our existing human image classifier, autonomously replicating the data analysis in Ref. [35]. In addition to developing a detection and positioning tool, we established a dataset of over 6000 labeled experimental images of BECs with and without solitonic excitations; this dataset is available via the National Institute of Standards and Technology (NIST) Science Data Portal [37] and at data.gov.

The remainder of this paper is organized as follows: in section 2, we illustrate the workflow of the soliton detector and its preparation process. Then in section 3, we demonstrate the system, quantify its performance, and discuss the quality of the labeled dataset. Finally in section 4, we conclude and discuss possible future directions.

2. Soliton detection and position system

In this section we describe our fully automated method of soliton detection and positioning in images of BECs. Our four-step protocol, detailed in the following sections and depicted in figure 1, is outlined as follows.

Step 1: Measurement. The measurement consists of three raw images that are combined to produce a single image of the atomic density distribution.

Step 2: Data preprocessing. As shown in figure 1, the BEC is rotated with respect to the image frame orientation, and the region of interest where atoms are captured is a small fraction of the full image. To simplify soliton positioning, the data is first rotated to align the BEC orientation with the image frame and then cropped prior to the classification step.

Step 3: Image classification. The pre-trained ConvNet classifier determines whether a lone soliton is present in a given image. If so, step four is executed, otherwise the image analysis terminates.

Step 4: Soliton positioning. The soliton position with respect to the BEC center is determined using a least-squares fit based on a one-dimensional (1D) model function.

2.1. Experimental setup and measurement

In our experiments, solitons are created and propagate the nonlinear media of a ^{87}Rb atomic BEC. We create BECs using well-established techniques for cooling and trapping atoms [38], allowing us to obtain $N = 2.4(2) \times 10^5$ atom³ condensates in a time-averaged crossed optical dipole trap. Since solitons are only stable in quasi-1D systems [39], i.e., resulting from highly anisotropic trapping geometries, our potential is elongated, with trapping frequencies $[\omega_x, \omega_y, \omega_z] = 2\pi \times [9.1(1), 153(1), 94.5(6)]$ Hz.

We launch solitons using our recently developed ‘improved’ protocol, that simultaneously engineers the density and phase of the BEC wave function [35]. By contrast with the ‘standard’ protocol that only modifies

³ We use a notation value(uncertainty) to express uncertainties, for example 1.5(6) cm would be interpreted as (1.5 ± 0.6) cm. All uncertainties herein reflect the uncorrelated combination of single-sigma statistical and systematic uncertainties.

the BEC phase and can only create solitons within a small range of initial velocities, our protocol can create solitons with arbitrary initial velocity. The potentials for density engineering and phase imprinting are both generated by far-detuned laser light, spatially patterned by a digital micromirror device (DMD). Our protocol is summarized as follows: After the BEC is created, we reduce its local density by applying a repulsive dimple potential. Next, the DMD is reprogrammed to display a step function that illuminates only half of the BEC, imprinting the soliton's phase profile. To minimize creating additional density perturbations, the dimple potential is reapplied and its magnitude slowly ramped to zero. We note that in our data there are additional solitonic excitations that, while representing different physical states (e.g. kink solitons, solitonic vortices, soliton rings and so forth [40]), can result in similar image and we identify simply as solitons in our analysis.

After solitons are created, we let them oscillate in the harmonic trapping potential for a variable evolution time. For evolution times much less than the trap period, additional density excitations from the soliton imprinting process are present. We then turn off the trapping potential and let the BEC evolve for a 15 ms time of flight, before absorption imaging the resulting density distribution [41].

2.2. Data preprocessing

We established a dataset of over 6.2×10^3 images for ConvNet training; these images were taken from multiple experiments performed in a single lab over a span of two months. The raw images were obtained with a 648×488 pixel camera (Point Grey FL3) with $5.6 \mu\text{m}$ square pixels, labeled by i and j . Including the $\approx 6\times$ magnification, each pixel has effective $0.93 \mu\text{m}$ size. The diffraction limit of the imaging system gives an optical resolution of $\approx 2.8 \mu\text{m}$ (roughly three pixels).

Absorption imaging combines three raw images into a single record of atomic density. In the first image $I_{i,j}^A$, a probe laser illuminates the BEC and the resulting intensity records the probe with the BEC's shadow. The second image $I_{i,j}^P$ records only the probe intensity, and the third image $I_{i,j}^{BG}$ is a dark frame containing any ambient background signal. The 2D column density

$$\sigma_0 n_{i,j} \approx -\ln \left[\frac{I_{i,j}^A - I_{i,j}^{BG}}{I_{i,j}^P - I_{i,j}^{BG}} \right] \quad (1)$$

can be derived from these images, where the resonant cross-section $\sigma_0 = 3\lambda^2/(2\pi)$ is derived from the wavelength λ of the probe laser. The dimensionless product $\sigma_0 n_{i,j}$ is of order 1 in our data, so we express density in terms of this product. Figure 1 shows an example of the probe beam with atoms and the resulting density in the 'raw data' and 'image classifier' frames, respectively.

In our raw data, the BEC occupies only a small region of the image, and the long axis of the BEC is rotated by about 43 degrees with respect to the camera. To facilitate the ConvNet training, the images are rotated to align the BEC with the image frame and cropped to discard the large fraction of the image that does not contain information about the BEC. Since the BEC's position and shape can vary for different realizations of the same experiment, we implement a fitting approach to determine the position and size of the BEC.

Next, we fit every image to a column-integrated 3D Thomas–Fermi distribution [42], giving the 2D distribution:

$$n_{i,j}^{\text{TF}} = n_0 \max \left\{ \left[1 - \left(\frac{i-i_0}{R_i} \right)^2 - \left(\frac{j-j_0}{R_j} \right)^2 \right], 0 \right\}^{3/2} + \delta n. \quad (2)$$

This function describes the density distribution of 3D BECs integrated along the imaging axis. We use six parameters to fit: the BEC center coordinates $[i_0, j_0]$; the peak 2D density n_0 ; the Thomas–Fermi radii $[R_i, R_j]$; and an offset δn from small changes in probe intensity between images.

Successful fitting requires acceptable initial guesses for all fit parameters. We obtained guesses for i_0 and j_0 by summing the image along the vertical and horizontal directions to obtain two 1D projections, from which we select the average position of the five largest values as the initial guesses. We took the largest value of the image as the guess for n_0 and used $[R_i, R_j] = [66, 55]$ pixels, based on the typical radii over the whole dataset. The guess for the offset δn is zero. The result of these fits are included in our released dataset.

We determined the 164×132 pixel extent of the cropping region by examining the radii $[R_i, R_j] = [66(5), 58(3)]$ obtained from fits to 6.2×10^3 images. We then centered the cropping region at $[i_0, j_0]$ as determined from fits of each image separately. The process was validated on an additional 10^4 images not included in our dataset. In the preprocessed images, dark solitons appear as vertically aligned density depletions and are easily visually identified (see top-left panel in figure 2(b)).

2.3. Labeling

Three independent human labelers labeled the preprocessed data, categorizing the images into three classes: 'no soliton', 'single soliton', and 'other excitations'. The 'no soliton' class contains images that unambiguously

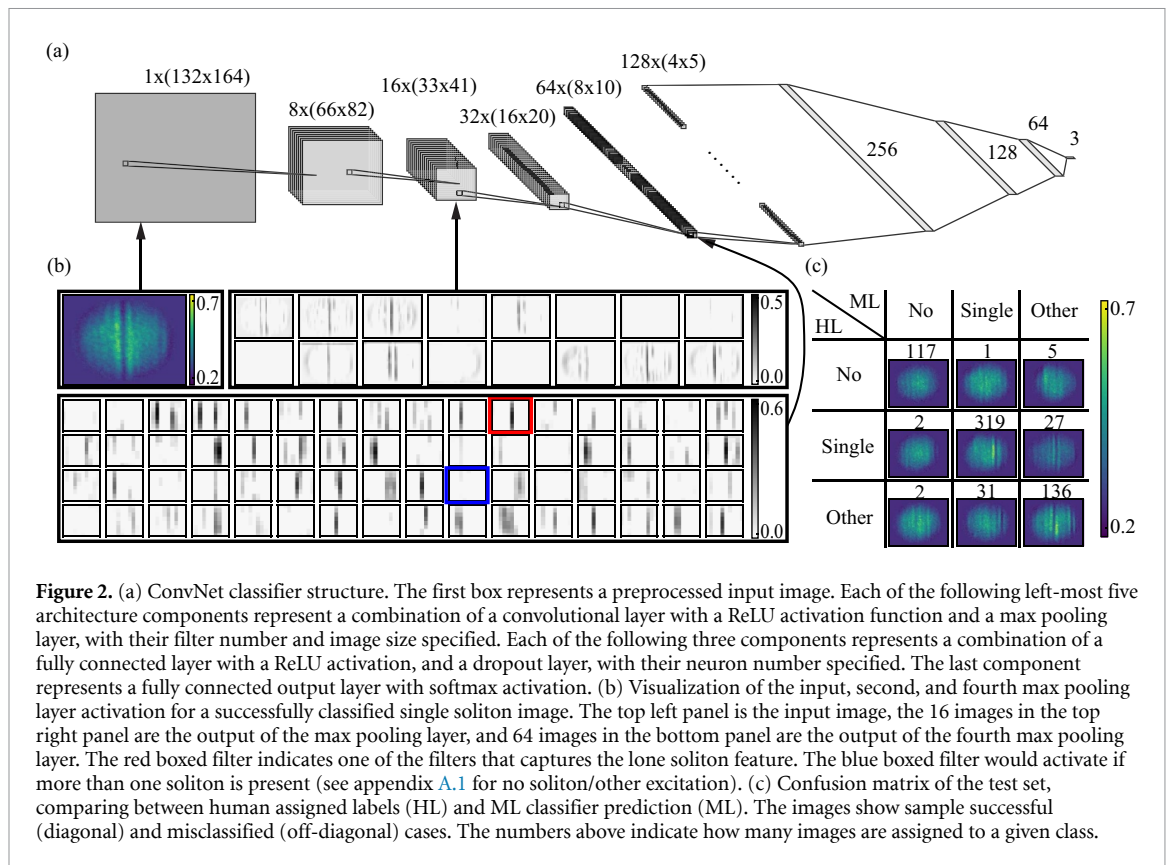


Table 1. Human labeling result. The first two columns (Full) show image counts and percentages of each class. The last two columns (3-agree) compare the counts and ratio in the all data of each class for the images with labels that humans initially agreed on.

Dataset Class	Full		3-agree	
	Count	Percentage [%]	Count	Agreement ratio [%]
No soliton	1237	19.8	1184	95.7
Single soliton	3468	55.4	3077	88.7
Other excitations	1552	24.8	1184	76.3
Total	6257	100.0	5445	87.0

contains no solitons; the ‘single soliton’ class describes images with one and only one soliton; and ‘other excitations’ class covers any image that can neither be interpreted as ‘no soliton’ nor ‘single soliton’. We did not include a separate ‘two soliton’ class in our demonstration because the small number of images with two solitons led to ineffective training.

The labeling process was carried out in eight batches, with each batch size limited by the attention span of the labelers. Once a given batch was completed, the resulting labels were compared and images with full agreement were set aside. The overall labeling agreement rate was 87% (table 1 shows a comparison of the labeling agreement for all three classes), consistent across all batches. The remaining images were further analyzed and discussed until an agreement was reached. The final distribution of images between classes is as follows: 19.8% in the no soliton class, 55.4% in the single soliton class, and 24.8% in the other excitations class. Figure 2(c) shows representative labeled images from each class. This labeled dataset was employed to train the ConvNet classifier and to test the positioning protocol.

2.4. Image classification

Our ConvNet classifier, shown in figure 2(a), consists of five convolutional layers. Each layer is followed by a rectified linear unit (ReLU) function defined as $f(x) = \max(0, x)$, then a max pooling layer⁴. The final max pooling layer is flattened and fully connected to a deep neural network with three hidden layers (256, 128, and 64 neurons, respectively) and an output layer (three neurons). Each hidden layer is followed by the ReLU

⁴ Max pooling is a form of non-linear down-sampling that converts the input ($km \times ln$) array, partitioned into a set of non-overlapping rectangles of equal ($k \times l$) size, into a smaller ($m \times n$) array with entries representing the maximum value of the corresponding sub-region.

Table 2. Classification performance summary for the best classifier when training with the full training dataset with performance measured using cross-validation from the training, when testing on the full test dataset, and when testing on a subset of the test dataset with labels that labelers initially agreed on.

	Cross-validation	Full test set	Labelers initially agreed subset
Accuracy [%]	89.6(5)	89.4	91.6
Weighted F1	0.896(6)	0.894	0.916
No soliton F1	0.938(10)	0.959	0.983
Single soliton F1	0.920(4)	0.913	0.935
Other excitations F1	0.806(6)	0.807	0.782

activation function, and to reduce overfitting, a dropout layer that randomly eliminates neural connections with a frequency of 0.5 during each training stage. The output vector $\xi = (\xi_1, \xi_2, \xi_3)$ is normalized by the softmax activation function, giving the final output probabilities $P_m(\xi) = \exp(\xi_m) / \sum_n \exp(\xi_n)$.

The labeled dataset was divided into two subsets: 640 images (10.2% of the dataset) were set aside as testing set, while the remaining 5617 images (89.8%) were used for training during the model architecture development. Since our training dataset is unbalanced, i.e. its different classes have a significantly different number of images, we balance it using augmentation techniques. We augment using three physically acceptable transformations: horizontal and vertical reflections, as well as a 180 degree rotation. All three transformations were applied to the no soliton and other excitations classes, increasing their size by a factor of four. For the single soliton class we used one randomly chosen transformation per image, doubling the size of this class. After augmentations, the size of the three classes has a 0.28 : 0.38 : 0.34 fractional distribution. To model a small rotation angle present in different realizations of our BEC, we randomly rotate images by an angle in the range ± 1 degree every time they are used during the training process. We applied an elliptical mask with radii $[R_i, R_j]$ to each image, eliminating all technical noise outside the BEC, to accelerate the training process⁵. Lastly, we preconditioned the data to have a range suitable for ConvNet input by uniformly scaling the image-values to the $[0, 1]$ range.

Since our testing dataset remains unbalanced, we assess the performance of trained models using the weighted F1 score [43]. When two models have similar weighted F1 scores, we first compare their accuracies as a tie-breaker, and if that fails we use the F1 score of the single soliton class⁶.

We used a semi-structured search through the model parameter space, and the resulting performance for varying hyperparameters is detailed in the appendix A.2. Once we determined the best performing model, we used randomly selected 95% of training set for the final training. Training terminated when the F1 score of the remaining 5% did not increase for five epochs. We took the model prior to these five non-improving epochs as our final trained model.

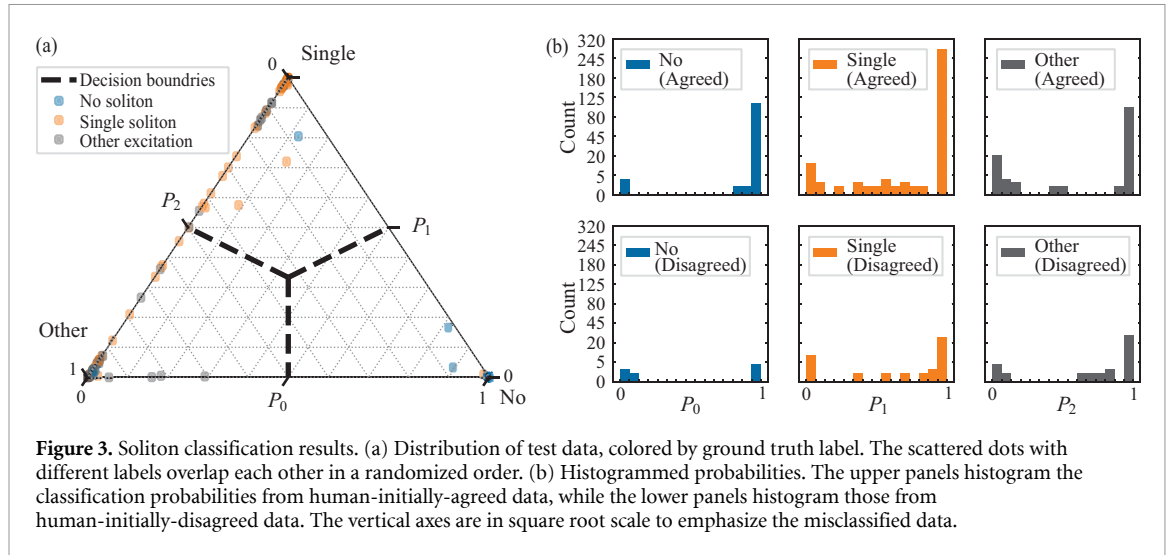
Figure 2(b) shows representative intermediate convolutional layers of the trained model, with a correctly classified single soliton as the input. We observe that some filters, such as the one marked with a red box, successfully capture the information of a single soliton (further examples are presented in appendix A.1).

Figure 2(c) and the second column of table 2 show the results of our final soliton classifier. In summary, our model has weighted $F_1 \approx 0.9$ and accuracy $\approx 90\%$, in excess of the 87.0% human agreement ratio. The most frequent classifier errors conflate images from the single soliton class and the other excitations class: 6.9% of the single soliton images is wrongly assigned to the other classes ($P_1 < 0.2$), and 4.3% has no clear assignment ($0.2 \leq P_1 < 0.8$).

Figure 3(b) shows that the classifier works very well for the no soliton and single soliton classes. The classifier performs better when tested against human-initially-agreed data than human-initially-disagreed data, suggesting that some disagreed upon images may be truly ambiguous (Also see the last column in table 2). In addition, we observe an anomalously large misclassification rate for human agreed data in the other excitations class, resulting from the human labelers use of this class when facing a dilemma. Furthermore, the wrongly classified data are distributed near the corners of figure 3(a), indicating a high degree of confidence in the misclassification.

⁵ In early training attempts, the classifier learned to separate the BEC from the background. Because the BEC resides within a well-defined ellipse, we accelerated convergence by applying the elliptical mask prior to training.

⁶ We use the F1 score of the single soliton class as the final tie-breaker, because we ultimately compare to single-soliton dynamics data in section 3.1.



2.5. Position regression

Once images containing only one soliton are identified, we locate the soliton position using a simple yet robust least-squares fitting procedure [43]. The first step consists of summing each 2D image along the j direction to obtain a 1D distribution $n_i = \sum_j n_{i,j}$. We fit the 1D distributions to the expected profile:

$$n_i^{1D} = n_0^{1D} \max \left\{ \left[1 - \left(\frac{i - i_0}{R_i} \right)^2 \right], 0 \right\}^2 + \delta n^{1D}, \quad (3)$$

that is, equation (2) integrated along the j direction. The initial guess for n_0^{1D} was the max of the integrated distribution, and the remaining guesses were taken from the successful 2D fit. We subtract the fit from the integrated 1D profile to obtain the residuals $\Delta_i = n_i - n_i^{1D}$. Ideally, this procedure would result in a flat background containing a single dip, associated with the soliton, which we identified using a Gaussian fit⁷. We use the minimum of Δ_i as the initial guess for the Gaussian amplitude, the minimum position as the initial center, 3 pixels for the width, and zero for the offset. This fit yielded the soliton width, amplitude and position.

3. Results

3.1. Soliton detector

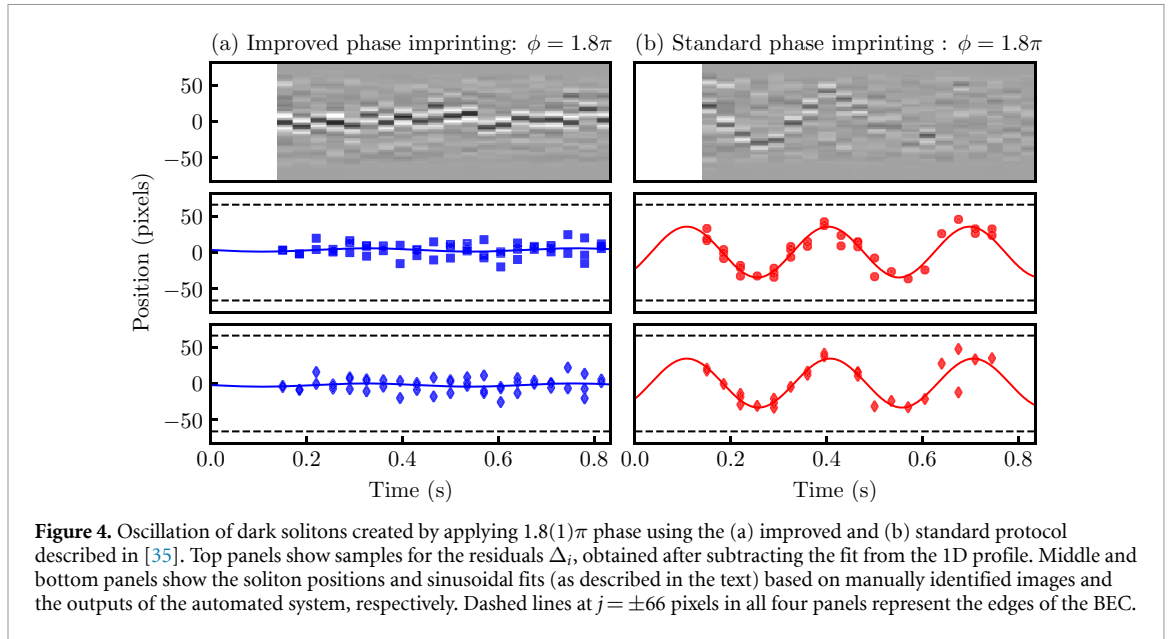
To test the performance of the fully automated soliton detection and positioning system, we use two sets of images containing oscillating dark solitons⁸ that were launched using the standard and improved protocols described in section 2.1, with 60 and 59 images, respectively.

In the first test, we used the improved-protocol data-set, with representative summed data n_i presented in the top panel of figure 4(a). As the solitons in these images are well pronounced, we expected the ConvNet will easily classify them. Out of 59 images, 52 were classified as single soliton and the remaining seven were classified as other excitations, in agreement with a human labeler. Solitons were then located in the first group by the positioning regressor (see figure 1). The middle and bottom rows in figure 4(a) plot the soliton position from manual and ConvNet identification, respectively. We fit $i(t) = A \sin(\omega t + \Phi) + i_0$ to the soliton position data, and we compare the fitted parameters with those obtained from our previous manual approach. As can be seen by comparing the middle and bottom rows of figure 4(a), the performance of the automated protocol is basically indistinguishable from the manual approach. The physical parameters from the ML classifier ($A = 2(2)$ pixels and $\omega/2\pi = 2.3(7)$ Hz) were within one standard deviation of those obtained for manual soliton identification ($A = 2(2)$ pixels and $\omega/2\pi = 2.3(6)$ Hz).

In the second test, we used images with solitons generated by the standard phase imprinting protocol. As can be seen in the top panel of figure 4(b), solitons in these images can be shallower than those in figure 4(a), making them potentially more difficult to distinguish from the no soliton and other excitations classes. Out of the 60 images in this test, 22 were classified by the ConvNet as no soliton, and 11 as other excitations, in

⁷ We found that the Gaussian gave the highest quality of fit as compared to other peaked functions.

⁸ These two sets of images contribute to the data published in figure 2 of [35], and are presented here in figure 4.



agreement with a human labeler. The remaining 27 were classified as a single soliton and were sent to the position regressor. The lower panels in figure 4(b) show soliton position as a function of evolution time, obtained from manual [35] and ConvNet identification, respectively. Since [35] compared the soliton oscillation amplitude resulting from the two imprinting protocols, the authors did not limit themselves to images with a single soliton. Rather, when more than one soliton was created, the authors identified all the solitons but tracked only that associated with a specific trajectory. Since the ConvNet classifier was trained to select images with single soliton excitations, the middle panel in figure 4(b) includes 12 more points than the bottom panel. Even with fewer data points, however, the parameters from the ML classifier ($A = 34(3)$ pixels and $\omega/2\pi = 3.34(9)$ Hz) were within one standard deviation of those obtained for manual soliton identification ($A = 35(2)$ pixels and $\omega/2\pi = 3.39(5)$ Hz).

The complete analysis resulting in both oscillation plots took under 148 s per series on a 2014 MacBook Pro. The expected performance relevant for in-situ operation is ≈ 2.4 s per image, a relatively small overhead on top of the measurement time (about 12 s). In many cases, however, the analysis of an image would take place during the acquisition of the next image.

3.2. Soliton dataset

As with all ML techniques, the availability of the training data is essential for good performance of the trained classifier. To assure the reliability of the assigned labels, the full dataset was independently labeled by three labelers, as described in section 2.2. Our full soliton image dataset consists of 6 257 labeled images. There are 1237, 3468, and 1552 images for no soliton, single soliton, and other excitations classes, respectively.

While for 5445 (87.0%) of the images the assigned labels were consistent between labelers, for the remaining 812 images (13.0%) there was a disagreement with at least one labeler. These images needed to be further discussed until an agreement was reached. As can be seen in table 1, the most challenging was distinguishing between images with single soliton and other excitations. This is likely due to the fact that the phase-imprinting method used to imprint solitons can also create other excitations that appear as density modulations or fringes in the BEC. Examples of such modulation can be seen in the off-diagonal images in figure 2(c). Additional discussion of the misclassified and mislabeled data can be found in appendix A.3.

Our dataset includes the full-frame raw images, the cropped and rotated images as used in this study, as well as the set of the fitted integrated 2D Thomas–Fermi distribution parameters. This dataset is sufficient to reproduce our results but also to test fitted alternative models with varying cropping size or image resolution [37].

4. Conclusion and outlook

In this manuscript, we present an automated dark soliton detection and positioning system that combines ML-based image classification with standard fitting techniques to track soliton dynamics in experimental images of BECs. We show that the system performs on par with more traditional approaches that rely on human input for soliton identification, creating the opportunity to study soliton dynamics in large datasets.

We also make available the first dataset of images from a dark soliton BEC experiment, which provides an opportunity for the data science community to develop more sophisticated analysis tools and to further understand nonlinear many-body physics.

The performance of the classifier, as measured by the weighted F1 score, leaves room for improvement. While tuning the hyperparameters allowed us to substantially improve the initial performance, additional data is necessary to push the limits. However, human labeling is not only time-consuming but, as the analysis of the misclassified images revealed, is also not always reliable. Other approaches, such as active learning ML [44], may be more suitable for this task. Such enlarged dataset, in turn, will enable refining the soliton classifier and perform model uncertainty quantification [45, 46], which currently is not accounted for. Together, these refinements may enable reliable in-situ deployment.

This study was preconditioned on the assumption of specific structure in the images, leading to our three classes. Enlarged dataset will enable employing unsupervised learning strategies [47] to possibly discover additional classes consistent with the data without presumptions. This unsupervised learning of soliton-data is a prototype for ML based discovery with cold-atom data in general.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.18434/mds2-2363>.

Acknowledgments

This work was partially supported by NIST and NSF through the Physics Frontier Center at the JQI. We appreciate conversations with Yuchen Yue and Justin Elenewski.

Appendix

A.1. Additional visualization of intermediate convolutional layers

In addition to figure 2(b) in the main text (visualizing the single soliton case), figure A1 shows the same intermediate layers for the correctly classified sample images from the (a) no soliton and (b) other excitations classes. In both figures, we highlight two filters: red box indicates a filter that activates for an image with a single soliton, while the blue box indicates a filter that activates for an image from the other excitations class. For an image from the no soliton class (figure A1(a)), neither highlighted filter is activated. This confirms that filters in our model are trained to properly detect and locate features that are characteristic of each class in a previously unknown image.

A.2. Model parameter tuning

We used a naive semi-structure search approach to optimize our model parameter. During parameter tuning, we used k -fold cross-validation to assess the generalizability of trained models. For each set of hyperparameters defining the overall model (e.g. kernel size and number of layers, both convolutional and hidden), the training set was split into $k = 5$ smaller sets ('folds'), four of which were used for training and one for validation. Once the model was trained using all five cross-validation combinations, the mean score was recorded and compared against scores from networks with other hyperparameters. We arrange the parameters by their importance: filter number of each convolutional layers, dense layer node number, optimizer, convolution kernel size, dropout rate and batch size. The parameters are optimized in this order and the history of hyperparameter tuning is shown in table A1.

A.3. Misclassified data

As discussed in section 2.1, in BEC experiments trapped condensates are often held for a certain period of time after phase imprinting. This is necessary to smooth out the various other excitations resulting from the phase imprinting process. However, by looking only at a single image (single holding time), all the information about the soliton dynamics is lost, and other excitations can be confused with shallow solitons. In the final 640 images in the test dataset, there are 68 misclassified images in one of six possible ways. As can be seen in figure 2(c), for our model, the most confusion comes from distinguishing between the single soliton and other excitations classes. Upon reviewing the 58 images misclassified in this way, we find that out of 27 images classified as other excitations only two clearly contain lone solitons. The remaining 25 images are confusing and thus should belong to the other excitations class. Interestingly, for this case, the classifier seems to be nearly as likely to be wrong as confused (see middle columns in figure 3(b)). All 31 images classified as

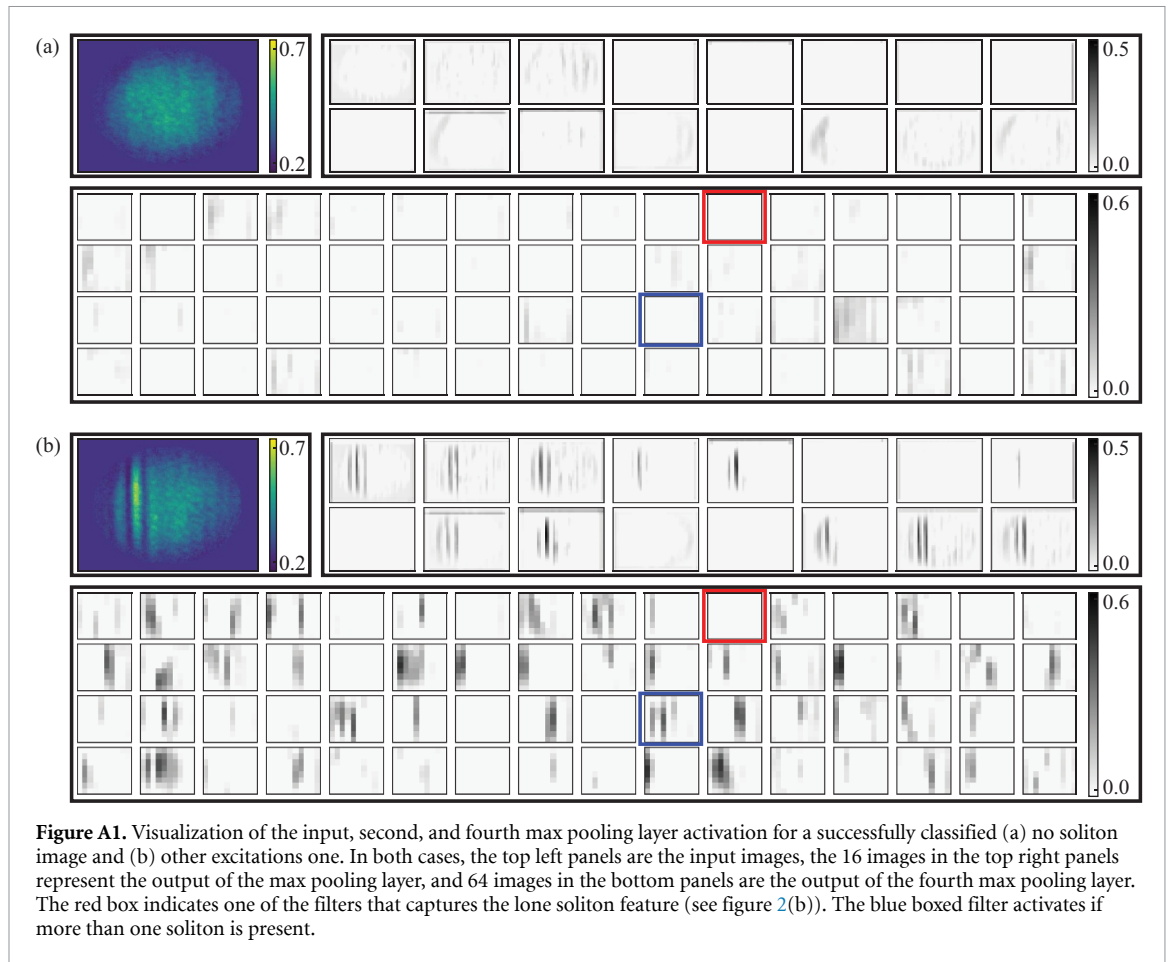


Figure A1. Visualization of the input, second, and fourth max pooling layer activation for a successfully classified (a) no soliton image and (b) other excitations one. In both cases, the top left panels are the input images, the 16 images in the top right panels represent the output of the max pooling layer, and 64 images in the bottom panels are the output of the fourth max pooling layer. The red box indicates one of the filters that captures the lone soliton feature (see figure 2(b)). The blue boxed filter activates if more than one soliton is present.

Table A1. Model parameter tuning history. For each model, we provide the number of filters for all convolutional layers (Filters), the number of nodes in the fully connected layers (Dense), as well as the kernel size of all convolutional layers (K), the dropout rate (D), the batch size (B), and the optimizer used of training (SGD: Stochastic gradient descent, SGD+M: SGD with moment, SGD+M+D: SGD+M with decay). The mean performance is averaged by 5-fold cross-validation on the training set. F1 score is weighted by three classes. Best model is highlighted. Parameters changed at each stage are highlighted in italic.

Filters	Dense	K	D	B	Optimizer	Weighted F1 [%]	Accuracy [%]	Binary F1 [%]
$8 \times 8 \times 8$	256×128	5	0.5	32	Adam	65(23)	74(15)	71(18)
$8 \times 16 \times 32$	256×128	5	0.5	32	Adam	56(23)	61(25)	63(19)
$8 \times 16 \times 32 \times 64$	256×128	5	0.5	32	Adam	67(24)	75(16)	72(20)
$8 \times 16 \times 32 \times 64 \times 128$	256×128	5	0.5	32	Adam	78(20)	82(14)	81(16)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.5	32	Adam	48(20)	62(13)	56(16)
$8 \times 16 \times 32 \times 64 \times 128$	256×64	5	0.5	32	Adam	58(25)	69(17)	65(20)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 64 \times 16$	5	0.5	32	Adam	61(19)	75(16)	69(17)
$8 \times 16 \times 32 \times 64 \times 128$	512×128	5	0.5	32	Adam	47(20)	54(22)	56(16)
$8 \times 16 \times 32 \times 64 \times 128$	$512 \times 128 \times 32$	5	0.5	32	Adam	67(24)	76(16)	72(20)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.5	32	<i>Adamax</i>	86(6)	89(1)	88(3)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.5	32	<i>SGD</i>	70(6)	87(2)	76(4)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.5	32	<i>SGD+M</i>	64(22)	74(15)	70(18)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.5	32	<i>SGD+M+D</i>	39(4)	60(9)	49(2)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.6	32	Adamax	77(5)	90(0)	85(3)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.7	32	Adamax	51(17)	68(15)	61(16)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	5	0.8	32	Adamax	44(13)	62(13)	54(12)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	3	0.5	32	Adamax	86(1)	90(1)	88(1)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	7	0.5	32	Adamax	88(1)	89(1)	90(1)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	9	0.5	32	Adamax	89(0)	89(0)	90(0)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	11	0.5	32	Adamax	78(20)	82(14)	82(17)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	9	0.5	16	Adamax	79(21)	83(13)	82(17)
$8 \times 16 \times 32 \times 64 \times 128$	$256 \times 128 \times 64$	9	0.5	8	Adamax	78(20)	82(13)	81(17)

single solitons are true misclassification. For these images, the classifier is confidently wrong (see last columns in figure 3(b)). We suspect that it might be possible to improve the classifier by training with less penalty for classifying an ambiguous image to other excitations class. This also suggests that active learning strategy might be better for training model and labeling data than relying on a dataset labeled by humans.

ORCID iDs

I B Spielman  <https://orcid.org/0000-0003-1421-8652>

Justyna P Zwolak  <https://orcid.org/0000-0002-2286-3208>

References

- [1] Aurisano A *et al* 2016 A convolutional neural network neutrino event classifier *J. Instrum.* **11** 09001
- [2] Acciarri R, Adams C, An R, Asaadi J and Auger M *et al* 2017 Convolutional neural networks applied to neutrino events in a liquid argon time projection chamber *J. Instrum.* **12** 03011
- [3] Kagan M 2020 Image-based jet analysis (arXiv:2012.09719)
- [4] Golovatiuk A, Giovanni D L and Ustyuzhanin A 2020 Deep learning for directional dark matter search *J. Phys.: Conf. Series* **1525** 012108
- [5] Khosa C K, Mars L, Richards J and Sanz V 2020 Convolutional neural networks for direct detection of dark matter *J. Phys. G: Nucl. Part. Phys.* **47** 095201
- [6] Kalantre S S, Zwolak J P, Ragole S, Xingyao W, Zimmerman N M, Stewart M D and Taylor J M 2017 Machine learning techniques for state recognition and auto-tuning in quantum dots *npj Quantum Inf.* **5** 1–10
- [7] Mills A R, Feldman M M, Monical C, Lewis P J, Larson K W, Mounce A M and Petta J R 2019 Computer-automated tuning procedures for semiconductor quantum dot arrays *Appl. Phys. Lett.* **115** 113501
- [8] Zwolak J P *et al* 2020 Autotuning of double-dot devices in situ with machine learning *Phys. Rev. Appl.* **13** 034075
- [9] Usman M, Wong Y Z, Hill C D and Hollenberg L C L 2020 Framework for atomic-level characterisation of quantum computer arrays by machine learning *npj Comput. Mater.* **6** 1–8
- [10] Cao Z, Dan Y, Xiong Z, Niu C, Xiang Li, Qian S and Jianjun H 2019 Convolutional neural networks for crystal material property prediction using hybrid orbital-field matrix and magpie descriptors *Crystals* **9** 191
- [11] Karamad M, Magar R, Shi Y, Siahrostami S, Gates I D and Farimani A B 2020 Orbital graph convolutional neural network for material property prediction *Phys. Rev. Mater.* **4** 093801
- [12] Gubernatis J E and Lookman T 2018 Machine learning in materials design and discovery: examples from the present and suggestions for the future *Phys. Rev. Mater.* **2** 120301
- [13] Duvenaud D K, Maclaurin D, Iparraguirre J, Bombarell R, Hirzel T, Aspuru-Guzik A and Adams R P 2015 Convolutional networks on graphs for learning molecular fingerprints *Adv. Neural Inf. Process. Syst.* **28** 2224–32 (<https://proceedings.neurips.cc/paper/2015/file/f9be311e65d81a9ad8150a60844bb94c-Paper.pdf>)
- [14] Butler K T, Davies D W, Cartwright H, Isayev O and Walsh A 2018 Machine learning for molecular and materials science *Nature* **559** 547–55
- [15] Winter R, Montanari F, Noé F and Clevert D-A 2019 Learning continuous and data-driven molecular descriptors by translating equivalent chemical representations *Chem. Sci.* **10** 1692–701
- [16] Rem B S, Käming N, Tarnowski M, Asteria L, Fläschner N, Becker C, Sengstock K and Weitenberg C 2019 Identifying quantum phase transitions using artificial neural networks on experimental data *Nat. Phys.* **15** 917–20
- [17] Ness G, Vainbaum A, Shkedrov C, Florshaim Y and Sagi Y 2020 Single-exposure absorption imaging of ultracold atoms using deep learning *Phys. Rev. Appl.* **14** 014011
- [18] Pilati S and Pieri P 2019 Supervised machine learning of ultracold atoms with speckle disorder *Sci. Rep.* **9** 5613
- [19] Metz F, Polo J, Weber N and Busch T 2021 Deep learning based quantum vortex detection in atomic Bose-Einstein condensates *Mach. Learn.: Sci. Technol.* (available at: <https://doi.org/10.1088/2632-2153/abea6a>)
- [20] Weller A, Ronzheimer J P, Gross C, Esteve J, Oberthaler M K, Frantzeskakis D J, Theocharis G and Kevrekidis P G 2008 Experimental observation of oscillating and interacting matter wave dark solitons *Phys. Rev. Lett.* **101** 130401
- [21] Frantzeskakis D J 2010 Dark solitons in atomic Bose-Einstein condensates: from theory to experiments *J. Phys. A: Math. Theor.* **43** 213001
- [22] Russel J S 1837 *Report of the Committee on Waves* 417–68 British reports VI plus plates 1–8
- [23] Osborne A R and Burch T L 1980 Internal solitons in the Andaman sea *Science* **208** 451–60
- [24] Lakshmanan M 2009 *Solitons, Tsunamis and Oceanographical Applications of* (New York: Springer) pp 8506–21
- [25] Burger S, Bongs K, Dettmer W, Ertmer W, Sengstock K, Sanpera A, Shlyapnikov G V and Lewenstein M 1999 Dark solitons in Bose-Einstein condensates *Phys. Rev. Lett.* **83** 5198–201
- [26] Denschlag J *et al* 2000 Generating solitons by phase engineering of a Bose-Einstein condensate *Science* **287** 97–101
- [27] Hasegawa A and Tappert F 1973 Transmission of stationary nonlinear optical pulses in dispersive dielectric fibers. ii. normal dispersion *Appl. Phys. Lett.* **23** 171–2
- [28] Mollenauer L F, Stolen R H and Gordon J P 1980 Experimental observation of picosecond pulse narrowing and solitons in optical fibers *Phys. Rev. Lett.* **45** 1095–8
- [29] Stasiewicz K, Shukla P K, Gustafsson G, Buchert S, Lavraud B, Thidé B and Klos Z 2003 Slow magnetosonic solitons detected by the cluster spacecraft *Phys. Rev. Lett.* **90** 085002
- [30] Hashizume Y 1985 Nonlinear pressure waves in a fluid-filled elastic tube *J. Phys. Soc. Japan* **54** 3305–12
- [31] Yomosa S 1987 Solitary waves in large blood vessels *J. Phys. Soc. Japan* **56** 506–20
- [32] Mollenauer L F and Gordon J P 2006 *Solitons in Optical Fibers* (Burlington, VA: Academic)
- [33] Hasegawa A 2000 Soliton-based optical communications: an overview *IEEE J. Sel. Top. Quantum Electron.* **6** 1161–72
- [34] Aycock L M, Hurst H M, Efimkin D K, Genkina D, Lu H-I, Galitski V M and Spielman I B 2017 Brownian motion of solitons in a Bose-Einstein condensate *Proc. Natl Acad. Sci.* **114** 2503–8
- [35] Fritsch A R, Mingwu L, Reid G H, Piñeiro A M and Spielman I B 2020 Creating solitons with controllable and near-zero velocity in Bose-Einstein condensates *Phys. Rev. A* **101** 053629

- [36] Rawat W and Wang Z 2017 Deep convolutional neural networks for image classification: a comprehensive review *Neural Comput.* **29** 2352–449
- [37] Database: data.nist.gov [Internet] Dark solitons in BECs dataset 2020 (available at: <https://doi.org/10.18434/mds2-2363>)
- [38] Lin Y-J, Perry A R, Compton R L, Spielman I B and Porto J V 2009 Rapid production of [sup 87]rb Bose-Einstein condensates in a combined magnetic and optical potential *Phys. Rev. A* **79** 63631
- [39] Muryshev A, Shlyapnikov G V, Ertmer W, Sengstock K and Lewenstein M 2002 Dynamics of dark solitons in elongated Bose-Einstein condensates *Phys. Rev. Lett.* **89** 110401
- [40] Muñoz Mateo A and Brand J 2015 Stability and dispersion relations of three-dimensional solitary waves in trapped Bose-Einstein condensates *New J. Phys.* **17** 125013
- [41] Ketterle W, Durfee D S and Stamper-Kurn D M 1999 Making, probing and understanding Bose-Einstein condensates *Bose-Einstein Condensation in Atomic Gases, Proc. Int. School of Physics 'Enrico Fermi', Course CXL* ed M Inguscio, S Stringari and C E Wieman (Amsterdam: IOS Press) pp 67–176
- [42] Castin Y and Dum R 1996 Bose-Einstein condensates in time dependent traps *Phys. Rev. Lett.* **77** 5315–19
- [43] Pedregosa F et al 2011 Scikit-learn: machine learning in Python *J. Mach. Learn. Res.* **12** 2825–30
- [44] Settles B 2009 Active learning literature survey *Technical Report* University of Wisconsin-Madison Department of Computer Sciences
- [45] Abdar M et al 2020 A review of uncertainty quantification in deep learning: techniques, applications and challenges (arXiv:2011.06225)
- [46] Thiebes S, Lins S and Sunyaev A 2020 Trustworthy artificial intelligence *Electron. Markets* **1–18**
- [47] Ji X, Henriques J F and Vedaldi A 2019 Invariant information clustering for unsupervised image classification and segmentation 2019 *IEEE/CVF Int. Conf. on Computer Vision (ICCV) (Seoul, Korea, Oct 27–Nov 2 2019)* (IEEE) 9864–73