

Combining machine learning with physics: A framework for tracking and sorting multiple dark solitons

Shangjie Guo  ¹, Sophia M. Koh  ^{2,3}, Amilson R. Fritsch  ¹, I. B. Spielman  ¹, and Justyna P. Zwolak  ^{3,*}

¹Joint Quantum Institute, National Institute of Standards and Technology, and University of Maryland, Gaithersburg, Maryland 20899, USA

²Department of Physics and Astronomy, Amherst College, Amherst, Massachusetts 01002, USA

³National Institute of Standards and Technology, Gaithersburg, Maryland 20899, USA



(Received 8 November 2021; accepted 5 May 2022; published 31 May 2022)

In ultracold-atom experiments, data often comes in the form of images which suffer information loss inherent in the techniques used to prepare and measure the system. This is particularly problematic when the processes of interest are complicated, such as interactions among excitations in Bose-Einstein condensates (BECs). In this paper, we describe a framework combining machine learning (ML) models with physics-based traditional analyses to identify and track multiple solitonic excitations in images of BECs. We use an ML-based object detector to locate the solitonic excitations and develop a physics-informed classifier to sort solitonic excitations into physically motivated subcategories. Lastly, we introduce a quality metric quantifying the likelihood that a specific feature is a longitudinal soliton. Our trained implementation of this framework, SOLDET, is publicly available as an open-source PYTHON package. SOLDET is broadly applicable to feature identification in cold-atom images when trained on a suitable user-provided dataset.

DOI: 10.1103/PhysRevResearch.4.023163

I. INTRODUCTION

Machine learning (ML) techniques promise improved data analysis and enhanced performance for today's quantum devices and technologies. Ultracold atomic gases are a nearly ideal system to deploy ML-driven analysis, where the automated exploration and interpretation of a very large dataset, in the form of images, can lead to scientific enhancements and experimental optimization [1] as well as new discoveries. Here we focus on the general problem of feature identification, a commonly recurring task in the analysis of such data, from locating vortices [2–4] or tracking solitons [5,6], identifying spin textures or magnetic domain walls [7–9] to locating topological singular points [10]. While data from these examples have been individually analyzed using task-specific algorithms (or even manual inspection), they are all feature identification problems that can be solved using a single ML-enhanced analysis framework. This paper introduces such a framework, and demonstrates its utility on the specific problem of identifying solitonic excitations in atomic Bose-Einstein condensates (BECs), as well as quantifying the quality of each identified feature.

Traditional statistical analysis using physics-based models, such as least-square fitting and hypotheses testing, have been go-to techniques for data analysis since the 1800's [11] and

remain widely applied in quantum cold-atom image analysis [12–14]. The outcome of physics-model-based algorithms and fits are intuitive, physically meaningful, and can help identify patterns present in the data; even fits based on more heuristic functions can have coefficients that are derived in obvious ways from the data. By contrast, ML methods work as "black boxes," making their operation difficult to interpret. Conventional statistical methods use fixed algorithms in conjunction with preconceived models for data reduction. Overfitting occurs when the number of fit parameters is comparable or larger than the number of independent data points. In this context, the process of training an ML tool essentially codesigns the fitting algorithm and the data model, as encoded by a large number of internal parameters. Training ML models is itself a fitting process that can be susceptible to overfitting, for example when the training dataset has too little variability or the ML model has too many internal parameters. ML involves a class of data-driven techniques that do not rely on preexisting models, but also add additional opportunities for overfitting that can make them less reliable on new data than conventional techniques.

Here, we describe the hybrid two-module feature identification framework shown in Fig. 1, that combines the flexibility of ML techniques with the intuition and robustness of conventional fitting methods. Furthermore the separate outputs of these two very different modules allow us to assess data quality by cross-validation. Hybrid approaches have been employed in other settings, for example for landslide prediction [15], medical image processing [16], and cyber attack detection [17].

The framework begins with a labeled dataset that is used to train the ML module and initialize the physics-based module. Before trusting either module, we independently validate each

*jpzwolak@nist.gov

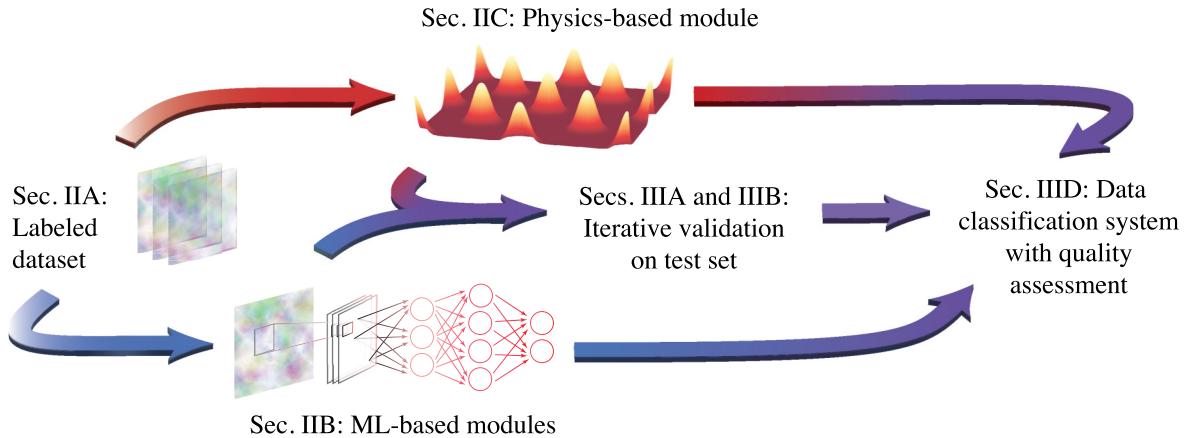


FIG. 1. Overview of the framework. The colored arrows link the preparation (Secs. IIA, IIB, and IIC), validation (Secs. IIIA and IIIB), and application (Sec. IIID) phases of the framework. The red path represents the preparation and implementation of the physics-based-approximation module of the framework. The blue path represents the ML modules.

module on a subset of the labeled data that was not used for training. Model redesign may be needed until satisfactory performance of each module is reached. We then combine both modules into an integrated system able to analyze new data.

We demonstrate the performance of our framework using data from atomic BECs, quintessential quantum systems. Quantum research with BECs, and cold-atom quantum gases more broadly, is multifaceted with examples ranging from realizing collective many-body physics [18] to creating today's most accurate atomic clocks [19]. In the vast majority of these experiments, data is acquired in the form of noisy images that typically have undergone evolution, such as a time of flight, before measurement. This often obfuscates the computation of the quantities of interest. Cold quantum gases therefore make an ideal testbed for our methodology that combines physically motivated, but heuristic, fitting functions with established computer vision techniques.

We focus on the specific problem of locating dark solitons (spatially compact excitations that manifest as reductions in the atomic density) as they move in BECs [13,20,21]. This allows us to leverage our established soliton dataset [22,23] to train and validate our framework; representative elements

of the dataset are shown in Fig. 2. These data consist of elliptical atom clouds (top row) where solitons appear as vertically aligned density depletions (bottom row). Not all vertically aligned density depletions are created equal: deep depletions mark the location of slowly moving kink solitons; shallow depletions are associated with rapidly moving kink solitons or “longitudinal” solitonic vortices (where the vortex core is aligned in the image plane); asymmetric depletions can result from “transverse” solitonic vortices [24] (where the vortex core is aligned perpendicularly to the image plane); and chains of stripes can result from highly excited phonon modes. Our framework is a tool that can automatically locate all the solitonic excitations in each image and distinguish between longitudinal solitons and transverse solitonic vortices. Here we introduce the term “longitudinal soliton” to include both kink solitons and longitudinal solitonic vortices.

Our ML module leverages and extends established computer vision techniques. Computer vision is a broad field with applications ranging from image classification to semantic segmentation and object detection [25]. Object detection refers to the capability of software systems to locate and identify objects in an image. Convolutional neural networks

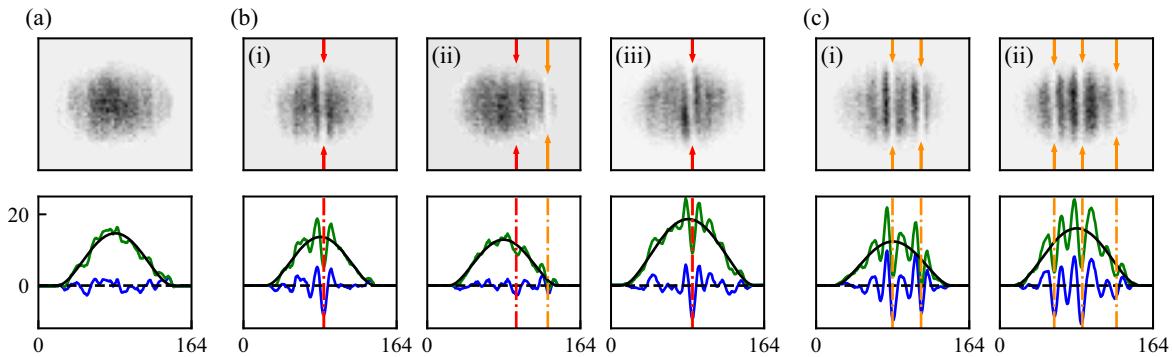


FIG. 2. Representative data. The top panels plot preprocessed images from our dataset and the bottom panels plot profiles: profile of full image (green), TF fits (black), and density fluctuations (blue). The red lines mark the location of the deepest depletion in the density fluctuations, while the orange lines mark the soliton locations found from our OD. (a) An element of the no-excitation class. (b) Three elements of the single-excitation class: (i) a single longitudinal soliton, (ii) an off-center longitudinal soliton, and (iii) a solitonic vortex. (c) Two representative elements of the other excitations class.

(CNNs) underlie solutions to all of these tasks, and unsurprisingly were employed in our previous work classifying soliton image data into three categories: no solitonic excitation, one solitonic excitation, and other excitations [22]. Our ML module goes beyond simple classification and uses a CNN-based object detector (OD) to provide the location of all candidate excitations in a given image.

By contrast our physics-based module employs a least-squares fit of an inverted and skewed Mexican-hat function to one-dimensional (1D) background-subtracted projections of soliton candidates (shown in bottom row in Fig. 2). We initialized this module using our previously labeled single soliton data and employ a Yeo-Johnson transformation [26] to produce a multivariate normal distribution yielding the likelihood that an unknown feature is a soliton.

This approach yielded three immediate benefits. First, a careful analysis of the coefficients from the physics-based module identified previously overlooked correlations that allow us to distinguish between some solitonic excitations (longitudinal solitons and transverse solitonic vortex [20,21,24,27]). Second, combining the results of the ML and fitting modules allowed us to automatically create a larger, more reliable dataset that includes fine-grained information such as the soliton position and type of excitation. This dataset is described in Ref. [28] and published in the NIST data repository [23]. Third, our hybrid framework was prepared solely from a training dataset whose images contain either zero or one solitonic excitation; however, it is performant on complex data containing multiple excitations.

The remainder of this paper is structured as follows: Section II introduces both modules and describes their training and initialization. Section III describes the validation of both modules and their performance on new data that include multiple solitonic excitations. In Sec. III E, we describe an open-source PYTHON reference implementation of our framework: SOLDET [29]. Lastly, in Sec. IV we conclude and discuss the potential applications of the framework as well as the possible future directions.

II. DATA AND MODULES

In addition to the recent success of ML methods [22,30,31], solitonic excitations have also been located and characterized using traditional fitting techniques. For example, Ref. [13] began with the background-removed atom density profiles (blue curves in Fig. 2) described in Sec. II A, then identified the deepest depletion (orange dashed line), and fit to a Gaussian function (a physically motivated, but heuristic choice) centered near the deepest depletion. This yielded physical information including soliton width, depth, and position. Unfortunately, this simple approach is failure prone, as for example in Fig. 2(b)(ii), where the deepest depletion is far from the actual soliton. Moreover, it detects only single solitonic features, making human intervention necessary when many excitations are present. Rather than finding the deepest minimum, our framework first uses an OD (described in Sec. II B) to provide an initial estimate of all solitonic excitation positions, and then uses a skewed Mexican-hat fit function (Sec. II C) that accurately describes their density profiles. The resulting fit coefficients serve two

purposes: qualitative likelihood assessment and fine-grained categorization.

A. Data

Our framework is trained and initialized using a revised dataset consisting of about 5.5×10^3 manually labeled experimental images of BECs with and without solitonic excitations [23,28]. The experimental setup and preprocessing techniques are described in [13].

Figure 2 shows six selected sample images from the labeled dataset. The dataset includes labels for five classes: “no solitonic excitation,” images that do not contain any excitations; “single solitonic excitation,” images containing one solitonic excitation; “other excitations,” images not in the preceding classes (including those with multiple solitonic excitations, high degrees of noise, and those annotators could not agree on); “mislabeled”, data determined to be potentially mislabeled during the curation process; and “unlabeled,” images that have not been manually annotated. Additionally, for the single excitation class the dataset includes the horizontal position of excitations within BEC.

Figure 2(a) displays an image from the no excitation class, which lacks the pronounced stripes present in the remaining examples. In (b), we show three elements of the single excitation class, each containing a single dark vertical fringe: (b)(i) a longitudinal soliton; (b)(ii) an off-center single longitudinal soliton; and (b)(iii) a solitonic vortex. In (c), we show two elements of the other excitations class containing more than one vertical fringe.

Horizontal 1D profiles (bottom row of Fig. 2) also have features associated with vertically aligned solitonic excitations and are amenable to least-squares fitting. We obtain these profiles by first summing the pixel values vertically to compress two-dimensional (2D) images to 1D; this sum can be over all (green curves) or part (see Sec. II C 1) of the vertical extent of the image. We then fit a 1D Thomas-Fermi (TF) model

$$n^{\text{TF}}(i) = n_0 \max \left\{ \left[1 - \left(\frac{i-i_0}{R_0} \right)^2 \right], 0 \right\}^2 + \delta_n \quad (1)$$

to each summed 1D profile, where i is the horizontal pixel index, and n_0 , i_0 , R_0 , and δ_n are fitting parameters representing peak density, center position, TF radius, and an overall offset, respectively. This fit (black curves) serves as an overall background that we subtract from the 1D profiles, leaving behind the 1D density fluctuations (blue curves). The orange dashed lines represent the location of deepest depletion in the 1D fluctuations.

B. ML modules

Our previous dark soliton classifier [22] consisted of a CNN model that returned one of the three predefined classes: no solitonic excitation, single solitonic excitation, or other excitations. However, this detector did not locate the excitations. To compare with experimental data, we located the soliton by identifying the deepest depletion and fitting to a Gaussian, as described above. This algorithm has two limitations: (1) The soliton may not be the deepest depletion [as in Fig. 2(b)(ii)]; and (2) multiple solitons cannot be located [as in Fig. 2(c)]. Here we retain the CNN classifier to globally organize the

data, but inspired by a highly successful recent result using an OD to locate vortices in numerically simulated 2D BECs [30], we employ an OD to locate solitonic excitations in experimental images of highly elongated BECs.

The OD complements the CNN classifier in two ways: (1) it identifies soliton positions rather than classifying; and (2) even though it is trained with single-soliton data, it can locate multiple excitations in the same image. We employ a neural network based OD with six convolution layers and four max-pooling layers but no fully connected layers (see the Appendix for more details). The OD has an order of magnitude fewer trainable parameters than our previous CNN (7×10^4 versus $\sim 10^6$ parameters), accelerating the training process and making it lightweight to deploy. Because the OD simply requires a dataset with many representative instances of the object to be detected, it requires far less training data than the CNN classifier (which by design required substantial data from all considered classes).

In our data, the solitonic excitations are roughly four pixels in width. Since our images are 164 pixels wide, we designed our OD to aggregate the image into 41 spatial cells, each with two outputs in the range $[0, 1]$; the OD therefore returns a 41×2 array $\tilde{\mathbf{Y}}$. For our dataset this aggregation guarantees that each output cell can describe the state of at most one soliton. $\tilde{Y}_{\ell,1}$ is a probability estimate that cell ℓ contains a soliton, and $\tilde{Y}_{\ell,2}$ is the fractional position of the soliton center within that cell, where 0 or 1 correspond to the left or right edge of the cell, respectively. The OD considers any cell with $\tilde{Y}_{\ell,1} > 0.5$ as containing an excitation, and then obtains its position from $\tilde{Y}_{\ell,2}$.

When comparing to the training dataset with labels denoted by \mathbf{Y} , we use the cost function [30]

$$F = \sum_{\ell=1}^{41} \begin{cases} -w_1 \log(\tilde{Y}_{\ell,1}) + w_2(Y_{\ell,2} - \tilde{Y}_{\ell,2})^2, & \text{if } Y_{\ell,1} = 1 \\ -\log(1 - \tilde{Y}_{\ell,1}), & \text{if } Y_{\ell,1} = 0 \end{cases} \quad (2)$$

for each training image, where the label $Y_{\ell,1}$ identifying the presence of an excitation in a cell is fully confident, i.e., either 0 or 1. The coefficients w_1, w_2 are hyperparameters controlling the relative importance of each term. The logarithmic terms increase the cost function when the OD misidentifies solitons, while the quadratic term contributes when a soliton is mislocated within a cell. Our training set uses images with at most one soliton, so cells with $Y_{\ell,1} = 1$ are much less frequent than those with $Y_{\ell,1} = 0$; as a result we expect $w_1, w_2 \gg 1$ to give similar overall weight to the three terms in Eq. (2). We train the OD by minimizing the cost function summed over all training images, updating the predicted OD values $\tilde{\mathbf{Y}}$ in each iteration. Because the cell size is comparable to the soliton size, a single soliton can span two cells. To prevent double counting, we merge detections occurring in adjacent cells and take the position to be their average.

We deem the OD's detection successful if our training data contains a labeled soliton close to the detected one (within three pixels in our implementation). The two failure modes are failing to detect a solitonic excitation and reporting an excitation that is not present.

C. Physics-based modules

In this section, we introduce our physics-based module that uses constrained least-squares fitting to estimate soliton parameters, and following a Yeo-Johnson transformation [26], produces a quality estimate giving the likelihood of a given feature being solitonic.

We fit the Ricker wavelet [32], i.e., a Mexican-hat function

$$f(i) = \delta_n - n^{\text{TF}}(i_c)A \exp\left[-\frac{1}{2}\left(\frac{i-i_c}{\sigma}\right)^2\right] \times \left[1 - a\left(\frac{i-i_c}{\sigma}\right)^2 + b\left(\frac{i-i_c}{\sigma}\right)\right], \quad (3)$$

to the 1D density fluctuations described Sec. II A, where $n^{\text{TF}}(i_c)$ is evaluated with $\delta_n = 0$. The function takes six parameters: normalized logarithmic amplitude A , center position i_c , width σ , logarithmic symmetrical shoulder height a , asymmetrical shoulder height b , and an offset δ . When a and b are zero this function is a simple Gaussian, making a nonzero adds symmetric shoulders to the distribution, and b introduces an asymmetry. Our solitonic features are well described by this function; since our excitations manifest as density depletions, the second term in Eq. (3) is negative.

Our constrained least-squares fit requires initial guesses for all of these parameters. The guess for the center position i_c also provides the initial guess for A by setting it equal to the 1D density fluctuations evaluated at i_c . We found the initial values $\sigma = 4$, $a = 0.2$, $b = 0$, and $\delta = 0$ to lead to convergent fits across the whole dataset. In order to produce reliable fits we apply the following constraints: i_c must remain within three pixels from the initial guess, $10^{-13} < A < 10^4$, and $10^{-13} < a < 10^4$ to prevent numerical fitting errors.

1. Physics-informed excitation classifier

Many candidate solitonic excitations are not vertically symmetric as might be expected [see, e.g., Fig. 2(b)(iii)]. The location of the largest “shoulder” in the top half of the excitation is reversed with respect to the bottom half; in addition, the location of the minimum is slightly displaced going from the top half to the bottom. Inspired by these differences, we bisect each image into top and bottom halves (labeled by + and -, respectively) and separately apply the Mexican-hat fit to fluctuations in these data, giving vectors Θ^\pm . Using this observation, we develop a physics-informed excitation (PIE) classifier based on the single-soliton dataset and discover that correlations between these vectors allow for a more fine-grained excitation classification.

Figure 3 shows the distribution of parameters from a single-soliton dataset that were useful for classifying excitations. No meaningful correlations were found for parameters σ^\pm and a^\pm , thus these did not assist in classification. The markers in the top panel show the amplitude ratio $\rho_A = A^+/A^-$ versus the top-bottom position difference $\delta i_c = i_c^+ - i_c^-$, and show that they are not correlated. By contrast, the bottom panel shows that the asymmetric shoulder height difference $\delta b = b^+/\sigma^+ - b^-/\sigma^-$ is clearly anticorrelated with δi_c . Both panels are colored based on the cut-off points discussed in Sec. III B (see also Fig. 5).

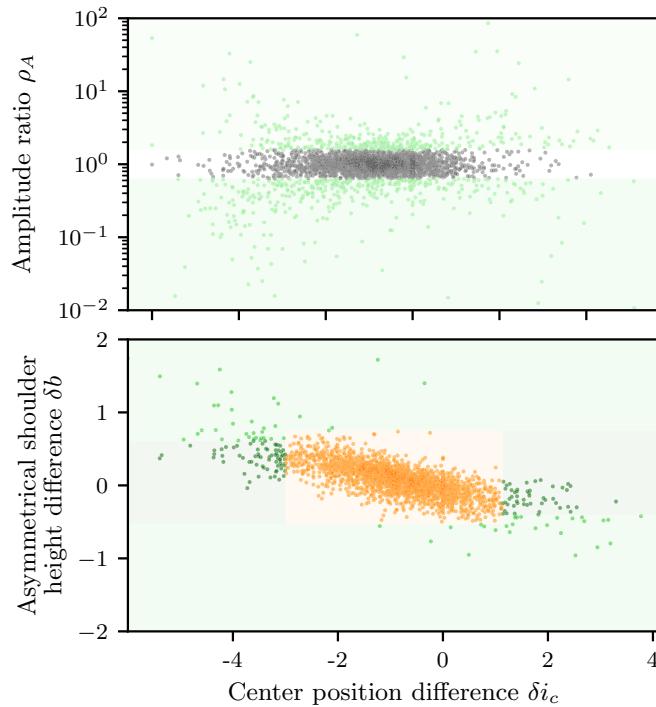


FIG. 3. Correlations between parameters implemented in PIE classifier. The top panel shows the distribution of center position difference versus the amplitude ratio (on a logarithmic scale). The bottom panel shows the correlation between the center position difference and the asymmetrical shoulder height difference for the gray points from the top panel. Both panels are colored based on the cut-off points discussed in Sec. III B

		0	1	2	3	4
Ground truth	0	1123	7	0	0	0
	1	20	3179	13	0	0
	'other'	82	352	575	24	3
		0	1	2	3	4
CNN classifier	0	1127	8	0	0	0
	1	26	3165	20	0	0
	'other'	72	365	568	24	3

FIG. 4. OD performance compared to ground truth (top), and the CNN classifier prediction (bottom). For ground truth and the CNN classifier, the ticks '0,' '1,' and 'other' represent no, single, and other excitation classes. For OD, ticks represent the total number of positive excitations within an image.

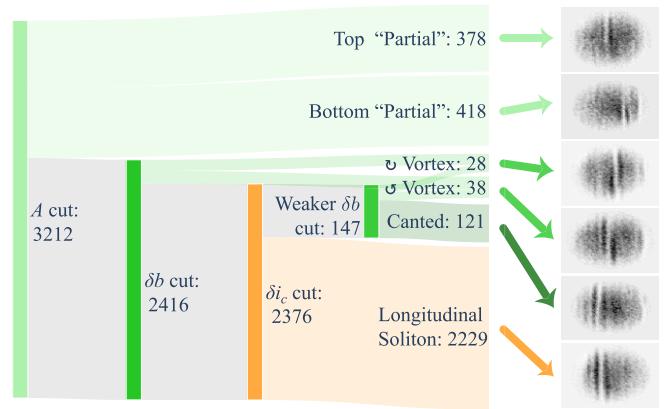


FIG. 5. The flow of the PIE classifier with example images for classification categories. Flow pathways and nodes are square-root scaled.

This distribution and its correlation guide the classification rules described in Sec. III B, yielding a PIE classifier based on cutoffs defined by human examination of the data.

2. Quality estimation

Here we describe a quality estimate that a candidate excitation in an image is solitonic. We derive the likelihood that a vector of fit outcomes $\Theta = [A, i_c, \sigma, a, b]$ is drawn from a $k = 5$ dimensional prior distribution spanning the set of representative solitonic excitations [33]. Ideally this distribution would be an uncorrelated multivariate normal distribution, but it is not. As a result, we developed the following procedure to bring the distribution into this desired form.

We first fit a Yeo-Johnson power transformation [26] to each separate parameter distribution (having summed the five-dimensional distribution along the remaining parameters) to transform them into independent zero-mean 1D Gaussian distributions with unit variance. Note that this treatment cannot transform the parameter distributions into perfect Gaussians; nevertheless, each resulting distribution is balanced, contains a single peak, and has long tails. The covariance matrix Σ_k is uncorrelated after this treatment and the distribution is qualitatively Gaussian in shape.

To calculate the quality estimate for a candidate excitation detected in an image, we

(1) fit the subtracted background 1D profile to Mexican-hat function given in Eq. (3) to obtain Θ ;

(2) use the established power transformation on Θ to obtain Θ' ; and

(3) return the quality estimate: $M(\Theta') = 1 - \chi_k^2[D^2(\Theta')]$, the likelihood between 0 and 1 that the excitation is solitonic.

The chi-squared cumulative distribution function $\chi_k^2(p)$ relates the Mahalanobis distance [34] $D^2(\Theta') = \Theta'^\dagger \Sigma_k^{-1} \Theta'$ to the likelihood that an outcome was drawn from the specified distribution [35]. $D(\Theta')$ is unbounded above and decreases to zero as Θ' approaches $\langle \Theta' \rangle$, the average over the prior distribution.

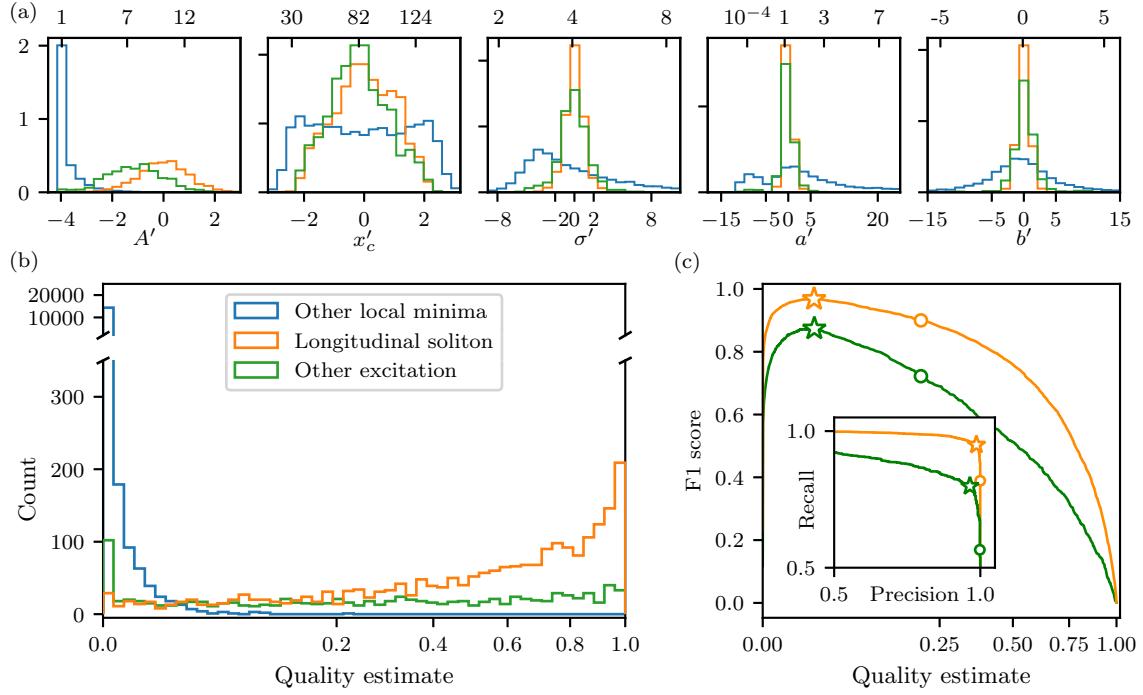


FIG. 6. Quality estimate performance on no excitation and single excitation classes. In all cases we use the following color scheme: longitudinal solitons (orange), all other solitonic excitations (green), and all nonsolitonic local minima (blue). (a) Power transformed fit coefficient distributions, with untransformed variables labeled on the top axis. (b) Distribution of quality estimate. (c) Performance of quality estimate quantified by F1 score; the stars indicate the optimal F1 value and the circles mark the threshold we use for classification. The inset shows precision/recall curves.

III. RESULTS

A. ML modules

We train both the CNN classifier and the OD using the refined dataset with added soliton position labels (see Ref. [28]). The CNN classifier is trained using the full dataset while the OD training uses only the no solitonic excitation and single solitonic excitation classes. We assess the performance of both modules using five-fold cross-validation, that is using 80% of the data to train a given module and the remaining 20% to test it, and repeating the process five times to fully cover the dataset (see the Appendix for training details).

The results are summarized in the two cumulative confusion matrices plotted in Fig. 4. The top panel compares the outcome of the OD to the initial labels, showing near perfect delineation between the no excitations and single excitations classes. However, the OD further subdivides the other excitations class, counting anywhere from zero to four candidate solitonic excitations within it. This results from the existence of excitations in this class that are not solitonic, as well as the possibility of having multiple solitons in the same image. The analogous comparison to CNN classification labels in the bottom panel is nearly indistinguishable from the one presented in the top panel, evidencing the quality of the CNN predictions.

Together, these ML tools effectively classify these data and locate excitations; however, they do not provide any fine-grained information on the nature or the quality of the identified excitations. This is addressed in the following sections.

B. PIE classifier

The PIE classifier operates by applying a sequence of “cuts” driven by different combinations of the top-bottom fit outcomes Θ^\pm . The exact parameter values described below are arrived at manually by exploring the data accepted and rejected by the cut to minimize the number of false-positive longitudinal soliton identifications.

The following cuts are applied sequentially, and the PIE classifier stops as soon as a classification is assigned.

A cut: The amplitude parameters A^\pm and their ratio ρ_A allow us to identify excitations that do not span the whole cloud. Data with $\rho_A > 1.57$ are classified as “top partial excitation” and those with $1/\rho_A > 1.57$ are classified as “bottom partial excitation.” This threshold identifies large fractional jumps in depth between the top and bottom that likely are off-axis solitonic vortices. Applying A cuts first is important because partial excitations interfere with the subsequent steps.

δb cut: Figure 2(b)(iii) illustrates a case with large shoulder height difference δb ; Ref. [27] showed that such data result from solitonic vortices. As a result, we classify data with $\delta b > 0.75$ as “counterclockwise solitonic vortex” and $\delta b < -0.53$ as “clockwise solitonic vortex.”

δi_c cut: Since longitudinal solitons have a vertically aligned density depletion [36], we classify data with $-3.0 < \delta i_c < 1.14$ as “longitudinal soliton.”

Weaker δb cut: Figure 3 shows that differences δi_c and $\delta b = b^+/\sigma^+ - b^-/\sigma^-$ are anticorrelated, indicating that asymmetries in position and shoulder height are related. A closer look at Fig. 2(b)(iii) indicates that it is such a case, with $\delta i_c < 0$ and $\delta b > 0$. We therefore add images with $\delta i_c < -3.0$

and $\delta b > 0.61$ to the counterclockwise solitonic vortex class and those with $\delta i_c > 1.14$ and $\delta b < -0.41$ to the clockwise solitonic vortex class.

Other data: The remaining images have $\delta i_c \neq 0$ but $\delta b \approx 0$ are labeled as “canted excitations,” likely kink solitons in the process of decay.

The flow chart in Fig. 5 shows the application of this classifier to a single-soliton dataset. We found that of the initial 3 212 images, about 1/3 failed a cut and were rejected as longitudinal soliton candidates.

This classification was also used in the preparation of Ref. [28] in which we present a refined soliton dataset, which includes improved single longitudinal soliton labels. The cuts above are fairly aggressive to avoid false positives in the longitudinal soliton classification. This implies possible misclassification in the other categories in order to ensure a high quality longitudinal soliton subset and a reliability of the quality metric.

C. Quality estimator

The quality estimator is initialized on the subset of the single excitation class identified as longitudinal soliton using the PIE classifier. Figure 6(a) shows the power-transformed distribution of Mexican-hat fit coefficients Θ' , with nontransformed coordinates marked on the top axis for reference. As would be expected, the data from the initialization dataset (orange) are nearly normally distributed; interestingly, the remaining elements of the single excitation class (partial solitons, canted excitations, and solitonic vortices, as labeled by the PIE filter) collectively follow very similar distributions (green). By contrast, the coefficients from every local minimum [37] in the initialization set *except* solitonic excitations (blue curve) obey a qualitatively different distribution.

Using this initialization, we compare quality estimates M obtained from the single excitation class in Fig. 6(b). The orange data show M for longitudinal solitons, and as intended the majority of this data is associated with larger values of M . The green data for the remaining solitonic excitations are nearly uniformly distributed, and the nonsoliton minima (blue) are highly concentrated at small M . We note that the small peak in longitudinal soliton distribution near-zero M contains a negligible fraction of the longitudinal soliton dataset (about 1.3%). However, this peak is more pronounced for the remaining excitations, which is not surprising because the power transform was initialized using longitudinal soliton data. These distributions demonstrate the ability of the quality estimator to discriminate between solitonic excitations and other features in the data, reinforcing the importance of the PIE filter for fine-grained classification.

We quantify the performance of the quality estimator in terms of the F1 scores plotted in Fig. 6(c), for longitudinal solitons (orange) and all other solitonic excitations (green). The F1 score for longitudinal solitons is maximized with a threshold of just $M = 0.02$ (stars): however, in practice we minimized false positives and assign a feature to be solitonic when $M > 0.2$ (circles). This choice gives only small change in the F1 score: however, it gives a marked increase in precision with only a small reduction in recall, as shown in the inset. The performance of the quality estimate on the other solitonic excitations, while far better than random, is subpar;

this reemphasizes the importance of the PIE classifier in our framework.

D. Application to other excitation and mislabeled data class

Here we discuss the performance of our SolDet framework applied to two classes of data from the dark soliton dataset: other excitations (1 036 images) and mislabeled data (879 images). These classes consist of images with multiple solitonic excitations, such as shown in Fig. 2(c), as well as confusing structures that made human annotation difficult. As such, they are an ideal test dataset since they defeated previous labeling attempts.

As a reminder, after the CNN classification step, the framework first uses the OD to locate all soliton candidates that are then sorted by the PIE classifier. Here, we focus only on features identified as longitudinal solitons. Figure 7(a) plots the frequency of transformed Mexican-hat fit outcomes Θ' , giving distributions that for both classes are qualitatively the same as those in Fig. 6(a) for the labeled single solitons. By contrast, histograms of the quality estimate for longitudinal solitons detected in these two classes [panel (b)] have important differences. For the other excitations class ($N_{\text{longitudinal}} = 877$, $N_{\text{images}} = 669$), the distribution is nearly uniform, with a potential increase for the higher quality estimates ($M > 0.4$). For the mislabeled data ($N_{\text{longitudinal}} = 415$, $N_{\text{images}} = 398$), on the other hand, the quality estimate distribution follows a trend consistent with that observed in Fig. 6(b).

To better understand this discrepancy it is important to consider more carefully the differences between the two classes. According to the OD module, nearly 78% of images in the other excitation class contains two or more excitations. While for excitation spaced apart within the BEC, as in Fig. 7(c)(i), the individual fits to Mexican hat do not affect one another, the contrary holds for excitation captured in close proximity, as shown in Fig. 7(c)(ii). Qualitative differences between these images are quantified by the quality estimate. The quality estimate for the two well separated excitations in image (i) is 0.74 and 0.86. In image (ii), in contrast, even though both excitations are reminiscent of a longitudinal soliton, they are assigned a low quality, with $M_{\text{(ii)}} = [0.00, 0.01]$ from left to right. This is likely because the overlap in the adjacent shoulders significantly affects the relative fits. Given that the majority of data in this class contains multiple excitations, the unusually high frequency of the low quality is to be expected.

The mislabeled class, on the other hand, consists of images determined to be potentially mislabeled during the manual annotation (see Ref. [28] for details about the data curation process). These include over 320 images that the annotators found confusing (but in which ODs consistently found exactly one candidate excitation); over 190 images removed during curation from the single excitation class; and about 30 images originally assigned to the no excitation class (but in which the ODs also consistently found exactly one candidate excitation). Unsurprisingly, in almost 83% of these images the OD module found only one excitation. Two representative images from this set are shown in Figs. 7(c)(iii) and 7(c)(iv), with $M_{\text{(iii)}} = [0.92, 0.02]$ and $M_{\text{(iv)}} = 0.82$. The distribution of nonlongitudinal soliton quality estimate, shown in the inset in Fig. 7(b), is consistent with that depicted in Fig. 6(b).

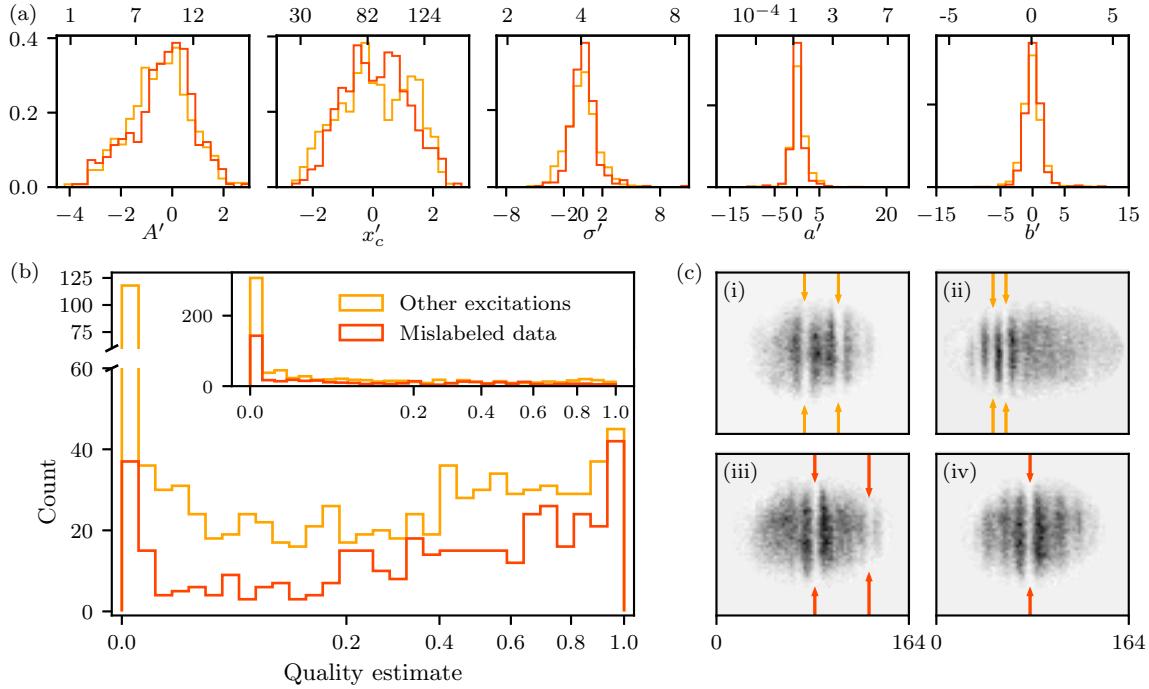


FIG. 7. Performance of quality estimate on other excitation (orange) and mislabeled (red) classes. (a) Power-transformed fit coefficient distributions, with untransformed variables labeled on the top axis. (b) Distribution of quality estimate of all longitudinal solitons. (c) Representative images from the other excitation [(i),(ii)] and mislabeled [(iii),(iv)] classes, with OD + PIE identified longitudinal soliton marked with arrows. The quality estimate for these excitations is as follows: $M_{(i)} = [0.74, 0.86]$, $M_{(ii)} = [0.00, 0.01]$, $M_{(iii)} = [0.92, 0.02]$ (all from left to right), and $M_{(iv)} = 0.82$.

The performance on these qualitatively different test sets emphasizes the power of SOLDET. By combining the CNN and OD modules, SOLDET autonomously and reliably locates multiple excitations within the BECs, which goes beyond the traditional state-of-the-art deepest-depletion-based approach. The PIE classifier enables further systematic validation that the desired type of excitation (here, longitudinal solitons) has been observed, which previously required visual inspection of each acquired image. Finally, the quality metric provides a quantitative assessment of the excitation quality, further reinforcing the classification reliability. Put together, these tools provide a robust and reliable analysis framework, capable of processing data significantly more complex than possible given the current traditional state-of-the-art approaches.

E. SOLDET: Open-source PYTHON package for solitonic excitation detection

In this section, we describe our software package SOLDET that integrates both the ML modules (CNN classifier and OD) with the fitting physics-based modules (PIE classifier and quality estimator), as we described in previous sections. The above discussion showed that the ML modules classify images effectively and can accurately locate one or many candidate solitons. The physics-based modules can sort these candidates into subclasses and provide a quality estimate for longitudinal soliton candidates. Therefore, the ML and physics-based modules contribute to the task of soliton detection in different ways, and the SOLDET infrastructure leverages their complementing strengths. We emphasize that soliton detection is one

of a larger class of feature identification in quantum gases and that SOLDET was designed to be broadly applicable.

The SOLDET distribution includes a CNN classifier, OD, PIE classifier, and quality estimator trained and initialized using the soliton dataset [23]. In addition, we provide training scripts to enable the ready application to user-defined data with custom preprocessors, ML models, fitting functions, and even the overall process flow.

Figure 8 illustrates a single use of SOLDET for the specific example of longitudinal soliton detection, where the individual blocks operate as follows:

Data processing: Preprocess raw data into a 164×132 image format that just encloses the elliptical atom clouds [22]. The preprocessing particulars are not generic and instead are specific to both our task as well as the experimental parameters.

CNN classifier: Apply a trained CNN classifier to processed data and yield labels no excitation, single excitation, or other excitations.

Object detection: Apply trained OD to processed data and yield a list of positions of solitonic excitations.

CNN:0 OR OD:0: If either the CNN classifier or OD finds no soliton, SOLDET terminates.

PIE classifier: The PIE classifier is applied to each solitonic excitation.

Quality estimator: The quality estimator is applied to each excitation identified as “longitudinal soliton” by the PIE classifier.

This algorithm is designed to be usable in a laboratory environment where one needs real-time identification, as well as for automated labeling of large datasets, as in Ref. [28].

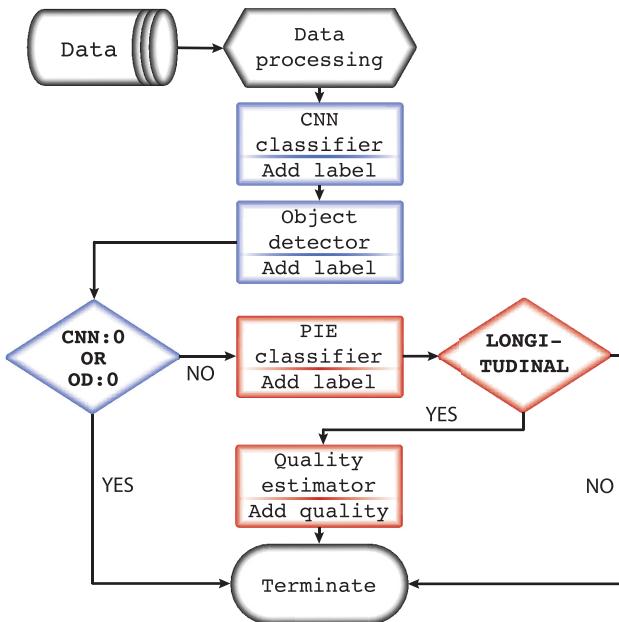


FIG. 8. The SOLDET flow chart. The black line follows the SOLDET dataflow and contains the labels added by each module (rectangles). Blue blocks represent ML modules; red blocks represent physics-based modules.

IV. DISCUSSION AND OUTLOOK

Here we described a framework that adds to the growing ML quantum science and technology toolkit, with additional recent developments including noise characterization [38,39]; quantum state detection [22,30,40–48]; parameter space exploration and optimization [49–54]; and quantum control [55,56]. Together, these results show that ML techniques can extract information from ambiguous data, efficiently search large parameter spaces, and optimally control quantum systems.

Our high level framework combines ML methods with physics-based analysis, providing an integrated platform for studying experimental data. Our implementation of this framework, SOLDET, currently targets the identification, classification, and tracking of features in image data generated by cold-atom experiments. We demonstrated its initialization and performance using a publicly available dark soliton dataset [28]. This investigation focused only on properties of individual images; however, the dataset also includes a label giving time elapsed since the excitation's were created. This opens the door for studies correlating system control parameters and the SOLDET labels.

While our initialization used only the no excitation and single excitation classes, SOLDET's feature detection successfully generalizes the learned patterns. This is confirmed by its performance on the other excitations and mislabeled classes that were not part of training, where the CNN classifier gave ambiguous results and human classifiers often disagreed. Going beyond simple classification tasks, SOLDET allowed us to identify unexpected structure in the data, enabling a fine-grained division of the single excitation class into physically relevant subclasses, including solitonic vortices and partial solitons.

Moreover, for the multiple excitations class, the distribution of the quality metric in Fig. 7 reveals a possible correlation between the quality metric and the excitations relative proximity. These observations illustrate the power of our combined framework as a data analysis tool for discovery.

An interesting application of SOLDET would be an off-line optimization of the experimental setup. Such optimization strategy, successfully implemented to, e.g., improve fabrication of quantum dot devices [1], requires an efficient analysis of large volumes of data to find the appropriate correlations in a high-dimensional parameter space. The ML toolbox described in our paper allows one to automatically locate multiple solitonic excitations in the same cloud and produces a fine classification that goes beyond longitudinal solitons. An analysis of the correlations between the various control parameter ranges used in our experiments and the resulting class of data (as determined by SOLDET) could enable a controlled generation of a desired number, type, and configurations of excitations, with SOLDET integrated online to provide real-time data analysis and control feedback. Another interesting extension of this work would be to train an OD on a dataset containing a single subclass found by the PIE classifier, e.g., longitudinal solitons, or solitonic vortices.

From the ML perspective, adding modules based on unsupervised [57], active learning [58], and synthetic data generation with generative models [59] may further enhance the performance of the SOLDET framework.

Going beyond solitonic excitations, the wakefield for sub- and supersonic impurities moving in atomic superfluids have characteristic patterns that could be identified by ML techniques [60–63]. This might be implemented using a template-based method such as used in the Laser Interferometer Gravitational-Wave Observatory (LIGO) where a large set of numerical simulations provide a library of patterns to correlate with the data [64]. This pattern matching is a form of object detection, and in our context a CNN based object detector could also be trained on such a template set. In this way, our methodology could be employed with a trained OD followed by a LIGO-like algorithm playing the role of our quality estimator and PIE classifier.

In the final analysis, SOLDET improves the data analysis pipeline for feature identification and classification problems in physically derived image data, but leaves the remainder of the scientific discovery process unchanged. For example, in our studies the PIE classifier module provided a fresh way to process data and enabled us to identify new patterns in the reduced data. The step beyond this is ML-driven discovery, where the identification of previously unknown patterns and physical reasoning are both implemented by ML. An emerging area of ML is the derivation of effective hydrodynamic equations of motion for biological, colloidal, and active fluids based on time-series data [65]. Owing to the complexity of full 3D simulations of nonzero temperature BECs, this data-driven approach could also be applied to create effective kinetic theory of solitons as well as the hydrodynamics of the underlying fluid.

ACKNOWLEDGMENTS

This work was partially supported by NIST and NSF through the Physics Frontier Center at the JQI. The views and

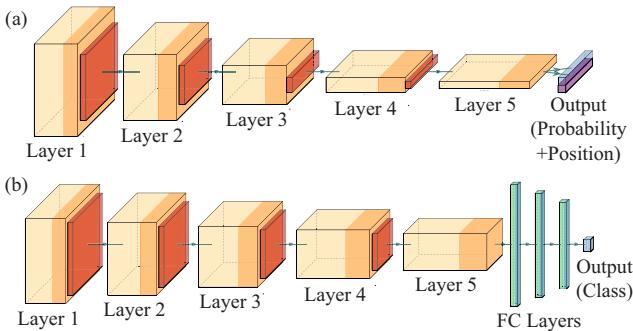


FIG. 9. Illustration of (a) OD and (b) CNN classifier neural network architectures. Yellow-orange boxes show convolutional layers while orange-red boxes show max-pooling layers. The horizontal lengths of boxes represent the number of filters and the other two dimensions represent the image sizes. The horizontal blue and purple rectangles in (a) denote output vectors. Each cell of the blue vector describes the probability that it contains a soliton and the purple vector contains the position of a soliton within the cell. And the vertical blue-green rectangles in (b) are three fully connected layers and the output layer. The lengths of edges are logarithmically scaled.

conclusions contained in this paper are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright noted herein.

APPENDIX: PARAMETERS OF MACHINE LEARNING MODELS

Both machine learning modules are built and trained using the TENSORFLOW (v.2.5.0) KERAS PYTHON API [66].

TABLE I. The OD architecture parameters. The top four rows are for the convolutional 2D layers and the three bottom rows are for max-pooling 2D layers.

Layer	1	2	3	4	5	Output
Filter	8	16	32	64	128	2
Kernel	5×5	5×5	5×5	1×5	1×5	1×5
Padding	Same	Same	Same	Same	Same	Same
Activation	ReLU	ReLU	ReLU	ReLU	ReLU	Sigmoid
Pool size	4×2	4×2	4×1	2×1	N/A	N/A
Strides	4×2	4×2	4×1	2×1	N/A	N/A
Padding	Valid	Valid	Same	Same	N/A	N/A

Figures 9(a) and 9(b) show the visualization of the network architecture for the OD and the CNN classifier, respectively. The model parameters of OD are presented in Table I. The model parameters for the CNN classifier are presented in the Appendix of Ref. [22].

As can be seen in Fig. 9, there are three main differences between the two architectures: (1) the OD outputs 41 local probabilities and positions while the CNN classifier only outputs one of three possible classes; (2) the CNN classifier contains three fully connected layers, which dramatically increase the number of trainable parameters, while OD does not; (3) the OD has asymmetric pool size and strides for vertical and horizontal directions, which are customized to the features in our dataset; the pool size and strides are symmetric for the CNN classifier. As a result, the OD has more than an order of magnitude fewer trainable parameters (7×10^4) than the CNN classifier (10^6).

- [1] A. B. Mei, I. Milosavljevic, A. L. Simpson, V. A. Smetanka, C. P. Feeney, S. M. Seguin, S. D. Ha, W. Ha, and M. D. Reed, Optimization of quantum-dot qubit fabrication via machine learning, *Appl. Phys. Lett.* **118**, 204001 (2021).
- [2] J. R. Abo-Shaeer, C. Raman, J. M. Vogels, and W. Ketterle, Observation of vortex lattices in Bose-Einstein condensates, *Science* **292**, 476 (2001).
- [3] V. Schweikhard, I. Coddington, P. Engels, V. P. Mogendorff, and E. A. Cornell, Rapidly Rotating Bose-Einstein Condensates in and near the Lowest Landau Level, *Phys. Rev. Lett.* **92**, 040404 (2004).
- [4] Y. J. Lin, R. L. Compton, K. Jiménez-García, J. V. Porto, and I. B. Spielman, Synthetic magnetic fields for ultracold neutral atoms, *Nature (London)* **462**, 628 (2009).
- [5] C. Becker, S. Stellmer, P. Soltan-Panahi, S. Dörscher, M. Baumert, E.-M. Richter, J. Kronjäger, K. Bongs, and K. Sengstock, Oscillations and interactions of dark and dark-bright solitons in Bose-Einstein condensates, *Nat. Phys.* **4**, 496 (2008).
- [6] L. M. Aycock, H. M. Hurst, D. K. Efimkin, D. Genkina, H.-I. Lu, V. M. Galitski, and I. B. Spielman, Brownian motion of solitons in a Bose-Einstein condensate, *Proc. Natl. Acad. Sci. USA* **114**, 2503 (2017).
- [7] L. E. Sadler, J. M. Higbie, S. R. Leslie, M. Vengalattore, and D. M. Stamper-Kurn, Spontaneous symmetry breaking in a quenched ferromagnetic spinor Bose-Einstein condensate, *Nature (London)* **443**, 312 (2006).
- [8] M. Vengalattore, J. Guzman, S. R. Leslie, F. Serwane, and D. M. Stamper-Kurn, Periodic spin textures in a degenerate $f = 1^{87}\text{Rb}$ spinor bose gas, *Phys. Rev. A* **81**, 053612 (2010).
- [9] S. De, D. L. Campbell, R. M. Price, A. Putra, B. M. Anderson, and I. B. Spielman, Quenched binary Bose-Einstein condensates: Spin-domain formation and coarsening, *Phys. Rev. A* **89**, 033631 (2014).
- [10] M. Tarnowski, M. Nuske, N. Fläschner, B. Rem, D. Vogel, L. Freystatzky, K. Sengstock, L. Mathey, and C. Weitenberg, Observation of Topological Bloch-State Defects and Their Merging Transition, *Phys. Rev. Lett.* **118**, 240403 (2017).
- [11] A. M. Legendre, *Nouvelles Méthodes pour la Détermination des Orbites des Comètes* (F. Didot, 1805).
- [12] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn, Making, probing and understanding Bose-Einstein condensates, in *Proceedings of the International School of Physics “Enrico Fermi”*, edited by M. Inguscio, S. Stringari, and C. E. Wieman (IOS Press, Amsterdam, 1999), pp. 67–176.

[13] A. R. Fritsch, M. Lu, G. H. Reid, A. M. Piñeiro, and I. B. Spielman, Creating solitons with controllable and near-zero velocity in Bose-Einstein condensates, *Phys. Rev. A* **101**, 053629 (2020).

[14] T. P. Purdy, K. E. Grutter, K. Srinivasan, and J. M. Taylor, Quantum correlations from a room-temperature optomechanical cavity, *Science* **356**, 1265 (2017).

[15] F. Huang, Z. Cao, J. Guo, S.-H. Jiang, S. Li, and Z. Guo, Comparisons of heuristic, general statistical and machine learning models for landslide susceptibility prediction and mapping, *Catena* **191**, 104580 (2020).

[16] S. Ghosh, M. R. Malgireddy, V. Chaudhary, and G. Dhillon, A new approach to automatic disc localization in clinical lumbar MRI: Combining machine learning with heuristics, in *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)* (IEEE, New York, 2012), pp. 114–117.

[17] J. Sakhnini, H. Karimipour, and A. Dehghantanha, Smart grid cyber attacks detection using supervised learning and heuristic feature selection, in *2019 IEEE 7th International Conference on Smart Energy Grid Engineering (SEGE)* (IEEE, New York, 2019), pp. 108–112.

[18] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms, *Nature (London)* **415**, 39 (2002).

[19] T. Bothwell, D. Kedar, E. Oelker, J. M. Robinson, S. L. Bromley, W. L. Tew, J. Ye, and C. J. Kennedy, JILA SrI optical lattice clock with uncertainty of 2.0×10^{-18} , *Metrologia* **56**, 065004 (2019).

[20] S. Burger, K. Bongs, S. Dettmer, W. Ertmer, K. Sengstock, A. Sanpera, G. V. Shlyapnikov, and M. Lewenstein, Dark Solitons in Bose-Einstein Condensates, *Phys. Rev. Lett.* **83**, 5198 (1999).

[21] J. Denschlag, J. E. Simsarian, D. L. Feder, C. Clark, L. A. Collins, J. Cubizolles, L. Deng, E. Hagley, K. Helmerson, W. P. Reinhardt, S. L. Rolston, B. I. Schneider, and W. D. Phillips, Generating solitons by phase engineering of a Bose-Einstein condensate, *Science* **287**, 97 (2000).

[22] S. Guo, A. R. Fritsch, C. Greenberg, I. B. Spielman, and J. P. Zwolak, Machine-learning enhanced dark soliton detection in Bose-Einstein condensates, *Mach. Learn.: Sci. Technol.* **2**, 035020 (2021).

[23] J. P. Zwolak, S. Guo, A. R. Fritsch, and I. B. Spielman, Dark solitons in BECs dataset. National Institute of Standards and Technology, <https://doi.org/10.18434/mds2-2363> (2021).

[24] A. M. Muñoz Mateo and J. Brand, Stability and dispersion relations of three-dimensional solitary waves in trapped Bose-Einstein condensates, *New J. Phys.* **17**, 125013 (2015).

[25] A. Voulodimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, Deep learning for computer vision: A brief review, *Comput. Intell. Neurosci.* **2018**, 7068349 (2018).

[26] I.-K. Yeo and R. A. Johnson, A new family of power transformations to improve normality or symmetry, *Biometrika* **87**, 954 (2000).

[27] S. Donadello, S. Serafini, M. Tylutki, L. P. Pitaevskii, F. Dalfovo, G. Lamporesi, and G. Ferrari, Observation of Solitonic Vortices in Bose-Einstein Condensates, *Phys. Rev. Lett.* **113**, 065302 (2014).

[28] A. R. Fritsch, S. M. Koh, S. Guo, I. B. Spielman, and J. P. Zwolak, Dark solitons in Bose-Einstein condensation: A dataset for many-body physics research, [arXiv:2205.09114](https://arxiv.org/abs/2205.09114).

[29] S. Guo, S. M. Koh, A. R. Fritsch, I. B. Spielman, and J. P. Zwolak, SolDet: Solitonic feature detection package (Version 1.0.0). GitHub, <https://doi.org/10.18434/mds2-2641> (2022).

[30] F. Metz, J. Polo, N. Weber, and T. Busch, Deep learning based quantum vortex detection in atomic Bose-Einstein condensates, *Mach. Learn.: Sci. Technol.* **2**, 035019 (2021).

[31] D. Leykam, I. Rondon, and D. G. Angelakis, Dark soliton detection using persistent homology, [arXiv:2107.14594](https://arxiv.org/abs/2107.14594).

[32] N. Ricker, Further developments in the wavelet theory of seismogram structure, *Seismol. Soc. Am., Bull.* **33**, 197 (1943).

[33] We found that δ was strongly correlated with the remaining five parameters and did not improve the quality estimate performance.

[34] P. C. Mahalanobis, On the generalized distance in statistics, National Institute of Science of India, 1936.

[35] This argument assumes no prior knowledge about the distribution of fit outcomes for structures that are not solitonic excitations.

[36] Our images are slightly rotated with respect to the vertical and horizontal axes leading to a small angle in our data.

[37] We require that the local minima be at least 7 pixels wide, i.e., a minimum at pixel i must obey $n_{i\pm j}^{1D} < n_{i\pm(j+1)}^{1D}$ for $j = 0, 1, 2$.

[38] R. Harper, S. T. Flammia, and J. J. Wallman, Efficient learning of quantum noise, *Nat. Phys.* **16**, 1184 (2020).

[39] J. Ziegler, T. McJunkin, E. S. Joseph, S. S. Kalantre, B. Harpt, D. E. Savage, M. G. Lagally, M. A. Eriksson, J. M. Taylor, and J. P. Zwolak, Toward Robust Autotuning of Noisy Quantum Dot Devices, *Phys. Rev. Applied* **17**, 024069 (2022).

[40] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nat. Phys.* **13**, 431 (2017).

[41] Y. Zhang and E.-A. Kim, Quantum Loop Topography for Machine Learning, *Phys. Rev. Lett.* **118**, 216401 (2017).

[42] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).

[43] J. Venderley, V. Khemani, and E.-A. Kim, Machine Learning Out-of-Equilibrium Phases of Matter, *Phys. Rev. Lett.* **120**, 257204 (2018).

[44] B. S. Rem, N. Käming, M. Tarnowski, L. Asteria, N. Fläschner, C. Becker, K. Sengstock, and C. Weitenberg, Identifying quantum phase transitions using artificial neural networks on experimental data, *Nat. Phys.* **15**, 917 (2019).

[45] C. Miles, A. Bohrdt, R. Wu, C. Chiu, M. Xu, G. Ji, M. Greiner, K. Q. Weinberger, E. Demler, and E.-A. Kim, Correlator convolutional neural networks: An interpretable architecture for image-like quantum matter data, *Nat. Commun.* **12**, 3905 (2021).

[46] P. Cha, P. Ginsparg, F. Wu, J. Carrasquilla, P. L. McMahon, and E.-A. Kim, Attention-based quantum tomography, *Mach. Learn.: Sci. Technol.* **3**, 01LT01 (2021).

[47] J. Venderley, M. Matty, K. Mallayya, M. Krogstad, J. Ruff, G. Pleiss, V. Kishore, D. Mandrus, D. Phelan, L. Poudel, A. G. Wilson, K. Weinberger, P. Upreti, M. R. Norman, S. Rosenkranz, R. Osborn, and E.-A. Kim, Harnessing Interpretable and unsupervised machine learning to address big data from modern X-ray diffraction, [arXiv:2008.03275](https://arxiv.org/abs/2008.03275).

[48] N. Maskara, M. Buchhold, M. Endres, and E. van Nieuwenburg, A learning algorithm with emergent scaling behavior for classifying phase transitions, [arXiv:2103.15855](https://arxiv.org/abs/2103.15855).

[49] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush, Fast machine-learning online optimization of ultra-cold-atom experiments, *Sci. Rep.* **6**, 25890 (2016).

[50] A. D. Tranter, H. J. Slatyer, M. R. Hush, A. C. Leung, J. L. Everett, K. V. Paul, P. Vernaz-Gris, P. K. Lam, B. C. Buchler, and G. T. Campbell, Multiparameter optimisation of a magneto-optical trap using deep learning, *Nat. Commun.* **9**, 4360 (2018).

[51] S. S. Kalantre, J. P. Zwolak, S. Ragole, X. Wu, N. M. Zimmerman, M. D. Stewart Jr., and J. M. Taylor, Machine learning techniques for state recognition and auto-tuning in quantum dots, *npj Quantum Inf.* **5**, 6 (2019).

[52] J. P. Zwolak, T. McJunkin, S. S. Kalantre, J. P. Dodson, E. R. MacQuarrie, D. E. Savage, M. G. Lagally, S. N. Coppersmith, M. A. Eriksson, and J. M. Taylor, Autotuning of Double-Dot Devices *In Situ* with Machine Learning, *Phys. Rev. Applied* **13**, 034075 (2020).

[53] A. J. Barker, H. Style, K. Luksch, S. Sunami, D. Garrick, F. Hill, C. J. Foot, and E. Bentine, Applying machine learning optimization methods to the production of a quantum gas, *Mach. Learn.: Sci. Technol.* **1**, 015007 (2020).

[54] H. Saito, Creation and manipulation of quantized vortices in Bose-Einstein condensates using reinforcement learning, *J. Phys. Soc. Jpn.* **89**, 074006 (2020).

[55] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. R. Carvalho, and M. J. Biercuk, Experimental deep reinforcement learning for error-robust gateset design on a superconducting quantum computer, *PRX Quantum* **2**, 040324 (2021).

[56] Mi.-Z. Ai, Y. Ding, Y. Ban, J. D. Martín-Guerrero, J. Casanova, J.-M. Cui, Y.-F. Huang, X. Chen, C.-F. Li, and G.-C. Guo, Experimentally realizing efficient quantum control with reinforcement learning, *Sci. China: Phys. Mech. Astron.* **65**, 250312 (2022).

[57] M. E. Celebi and K. Aydin, *Unsupervised Learning Algorithms* (Springer, New York, 2016).

[58] Li-Li Sun and Xi-Zhao Wang, A survey on active learning strategy, in *2010 International Conference on Machine Learning and Cybernetics* (IEEE, New York, 2010), Vol. 1, pp. 161–166.

[59] J. Gui, Z. Sun, Y. Wen, D. Tao, and J. Ye, A review on generative adversarial networks: Algorithms, theory, and applications, *arXiv:2001.06937*.

[60] I. Carusotto, S. X. Hu, L. A. Collins, and A. Smerzi, Bogoliubov-Čerenkov Radiation in a Bose-Einstein Condensate Flowing against an Obstacle, *Phys. Rev. Lett.* **97**, 260403 (2006).

[61] Yu. G. Gladush, G. A. El, A. Gammal, and A. M. Kamchatnov, Radiation of linear waves in the stationary flow of a Bose-Einstein condensate past an obstacle, *Phys. Rev. A* **75**, 033619 (2007).

[62] T.-L. Horng, S.-C. Gou, T.-C. Lin, G. A. El, A. P. Itin, and A. M. Kamchatnov, Stationary wave patterns generated by an impurity moving with supersonic velocity through a Bose-Einstein condensate, *Phys. Rev. A* **79**, 053619 (2009).

[63] A. M. Kamchatnov and L. P. Pitaevskii, Stabilization of Solitons Generated by a Supersonic Flow of Bose-Einstein Condensate Past an Obstacle, *Phys. Rev. Lett.* **100**, 160402 (2008).

[64] P. Ajith, S. Babak, Y. Chen, M. Hewitson, B. Krishnan, A. M. Sintes, J. T. Whelan, B. Brügmann, P. Diener, N. Dorband, J. Gonzalez, M. Hannam, S. Husa, D. Pollney, L. Rezzolla, L. Santamaría, U. Sperhake, and J. Thornburg, Template bank for gravitational waveforms from coalescing binary black holes: Nonspinning binaries, *Phys. Rev. D* **77**, 104017 (2008).

[65] R. Supekar, B. Song, A. Hestewell, G. P. T. Choi, A. Mietke, and J. Dunkel, Learning hydrodynamic equations for active matter from particle simulations and experiments, *arXiv:2101.06568*.

[66] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems, 2015 (software available from tensorflow.org).