



PAPER • OPEN ACCESS

CNN-Based vortex detection in atomic 2D Bose gases in the presence of a phononic background

To cite this article: M Sesodia *et al* 2025 *Mach. Learn.: Sci. Technol.* **6** 015067

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

You may also like

- [Stochastic resetting mitigates latent gradient bias of SGD from label noise](#)
Youngkyoung Bae, Yeongwoo Song and Hawoong Jeong
- [Improving early detection of gravitational waves from binary neutron stars using CNNs and FPGAs](#)
Ana Martins, Melissa Lopez, Gregory Baltus et al.
- [Towards virtual painting recolouring using vision transformer on x-ray fluorescence datacubes](#)
Alessandro Bombini, Fernando García-Avello Bofías, Francesca Giambi et al.



PAPER

OPEN ACCESS

RECEIVED
21 October 2024REVISED
28 February 2025ACCEPTED FOR PUBLICATION
12 March 2025PUBLISHED
21 March 2025

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.



CNN-Based vortex detection in atomic 2D Bose gases in the presence of a phononic background

M Sesodia , S Sunami , E Chang , E Rydow , C J Foot and A Beregi*

Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

* Author to whom any correspondence should be addressed.

E-mail: abel.beregi@physics.ox.ac.uk**Keywords:** machine learning, convolutional neural network, image analysis, Monte Carlo simulation, quantum vortices, superfluidity, Bose–Einstein condensate

Abstract

Quantum vortices play a crucial role in both equilibrium and dynamical phenomena in two-dimensional (2D) superfluid systems. Experimental detection of these excitations in 2D ultracold atomic gases typically involves examining density depletions in absorption images, however the presence of a significant phononic background renders the problem challenging, beyond the capability of simple algorithms or the human eye. Here, we utilize a convolutional neural network to detect vortices in the presence of strong long- and intermediate-length scale density modulations in finite-temperature 2D Bose gases. We train the model on datasets obtained from *ab initio* Monte Carlo simulations using the classical-field method for density and phase fluctuations, and Gross–Pitaevskii simulation of realistic expansion dynamics. We use the model to analyze experimental images and benchmark its performance by comparing the results to the matter-wave interferometric detection of vortices, confirming the observed scaling of vortex density across the Berezinskii–Kosterlitz–Thouless critical point. The combination of a relevant simulation pipeline with machine-learning methods is a key development towards the comprehensive understanding of complex vortex-phonon dynamics in out-of-equilibrium 2D quantum systems.

1. Introduction

Quantum vortices are topological excitations found in superfluids and are characterized by integer multiples of $\pm 2\pi$ phase winding around a vortex core where the density vanishes. They have been experimentally observed in a wide variety of systems including liquid helium [1, 2], type-II superconductors [3, 4], photon fields [5, 6], exciton–polariton superfluids [7, 8], and in ultracold atom systems with bosonic [9, 10] and fermionic [11–13] particles. The precise spatial and temporal control of the ultracold atom platform facilitates the study of numerous manifestations of vortex excitations as a consequence of rich static and dynamical many-body phenomena—for example, vortices have been shown to arrange into lattices in rotating Bose–Einstein condensates (BECs) [14, 15] and randomly form following a rapid quench into an ordered phase [16, 17]. An important class of vortex excitations emerges in systems within the 2D *XY* universality class to which ultracold 2D Bose gases belong [18], having characteristic paired and unpaired configurations across the superfluid transition [19] (see appendix A for a brief overview). The statistics and spatial structure of vortex excitations are one of the direct measures of the many-body states and their time–evolution, hence the precise determination of the number and location of vortices is critical for studying superfluid dynamics. Absorption imaging is typically used to obtain images of the density of ultracold atomic systems [20, 21], however the task of identifying vortices from these is often non-trivial due to their small spatial structure and misleading patterns arising from other excitations in the system. In some situations, it can be simple to count vortices by eye, e.g. in the ground state of 3D rotating BECs where these vortices form an ordered triangular lattice [14, 15]. However, in stationary states of finite-temperature 2D Bose gases the vortices do not form a regular pattern and the presence of phonons, which can result in similar density modulations after time-of-flight expansion [22, 23], makes vortex detection by the human

eye very inaccurate and impossible to use for quantitative analysis. Matter-wave interferometry addresses this challenge by taking advantage of the characteristic $\pm\pi$ phase discontinuity across (singly quantized) vortices [24]. These distinct features are easy to distinguish from other excitations which have a significantly longer characteristic length scale. This approach has accelerated 2D Bose gas research in recent years by providing a method to obtain the average vortex density from an ensemble of images [25–27], however this method is limited to measuring averages, missing the rich information available in single-shot detection of quantum vortex locations distributed in the plane of the 2D gas.

To address this constraint, here, we report on the development of a robust method for detecting vortex excitations in atomic 2D Bose gases by employing machine-learning techniques. In particular, we focus on the experimentally-relevant regime where thermally-activated vortices exist intermingled amongst other intrinsic excitations such as phonons which renders vortex detection by the human eye or simple algorithms challenging. Metz *et al* [28] showed that a convolutional neural network (CNN) trained on synthetic data of rotating 3D BECs can detect vortices in out-of-sample data, even in the presence of technical noise. Recently, Kim *et al* [29] applied the same approach to turbulent condensates, and accurately detected vortices in experimental data. Here, we focus on the challenging situation of detecting vortices arising from the Berezinskii–Kosterlitz–Thouless (BKT) phase transition, which are obscured by the presence of long- and intermediate-length scale density fluctuations by deploying a more comprehensive training and inference pipeline that is robust to both technical and intrinsic noise. We report automated vortex detection in experimental images well beyond the capability of counting by eye, and benchmark against the results obtained via matter-wave interferometry. The investigation of phase transitions at finite temperature is an important application for which our method can identify vortices in a non-uniform background density. More generally, we establish a framework for combining simulation with machine-learning methods which can be applied to a wide variety of systems.

The remainder of this paper is organized as follows. Section 2 describes how we numerically simulate 2D Bose gases to obtain experimentally-relevant training data. Section 3 outlines the design and operation of our CNN. Section 4 evaluates the model’s performance against validation data and examines its effectiveness at predicting vortex locations in experimental images. Finally, section 5 summarizes our work and provides an outlook for the present result.

2. Synthesizing training data

2.1. Numerical simulation of 2D Bose gases

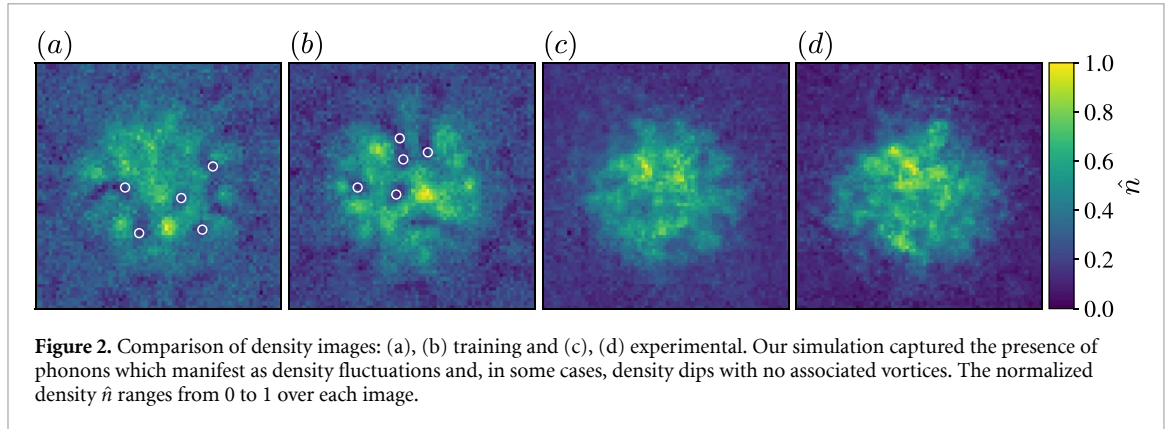
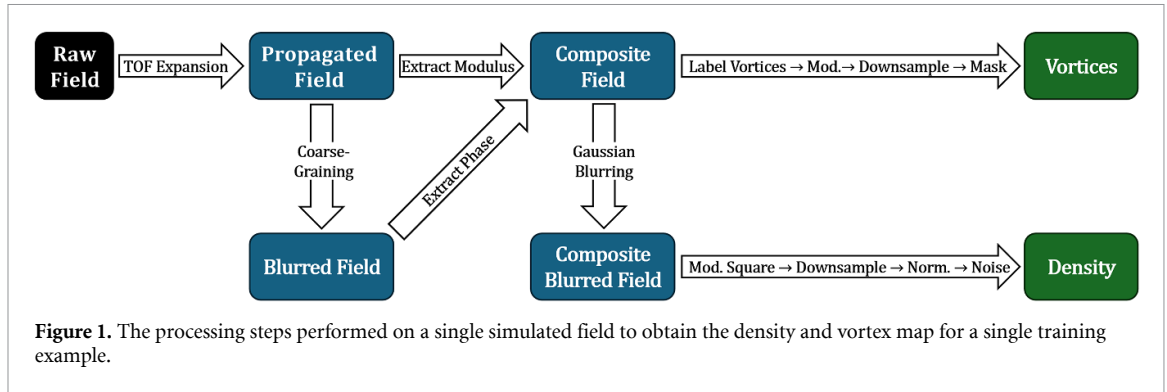
The physics of ultracold 2D Bose gases can be described by the Hamiltonian

$$\hat{H} = \int d\mathbf{r} \left[\frac{\hbar^2}{2m} \nabla \hat{\psi}^\dagger(\mathbf{r}) \cdot \nabla \hat{\psi}(\mathbf{r}) + \frac{g}{2} \hat{\psi}^\dagger(\mathbf{r}) \hat{\psi}^\dagger(\mathbf{r}) \hat{\psi}(\mathbf{r}) \hat{\psi}(\mathbf{r}) + V(\mathbf{r}) \hat{\psi}^\dagger(\mathbf{r}) \hat{\psi}(\mathbf{r}) \right], \quad (1)$$

where $\hat{\psi}(\mathbf{r})$ and $\hat{\psi}^\dagger(\mathbf{r})$ are annihilation and creation operators for bosons, respectively, at position vector \mathbf{r} in the x - y plane, $g = \tilde{g}\hbar^2/m$ is the quasi-2D interaction parameter, \tilde{g} is the dimensionless 2D interaction strength, and $V(\mathbf{r})$ is the external potential [30]. For the experimental parameters used the trapped gas exhibits near-homogeneous density, therefore we considered a round box potential of radius $22.5 \mu\text{m}$. This also allowed us to define the phase-space density of the gas as $\mathcal{D} = nh^2/(2\pi mk_B T) \propto n/T$, where n is the average 2D density, and T is the temperature (see appendix B for more experimental details).

To generate *in-situ* samples of the fluctuating Bose field for our system, we first used the classical field approximation to replace the field operators in (1) with complex numbers and their complex conjugates [31]. This approach accurately captured both vortex and phonon behavior, and has been previously used to support experimental observations of a wide range of superfluid phenomena in Bose gases [25, 32, 33]. We mapped the system onto a 280×280 square lattice with lattice constant $\delta l = 0.25 \mu\text{m}$ and evaluated the derivatives using first-order symmetric differences. The lattice constant was chosen smaller than the characteristic length scales of the system to ensure accurate representation of the continuum limit (see appendix C for this calculation).

We employed the Metropolis–Hastings algorithm [34, 35] to draw random states of a finite-temperature 2D Bose gas from the grand canonical ensemble. We initialized the complex field with zero real and imaginary components. At every iteration, we chose a random site and generated proposed updates to the real and imaginary parts of the on-site field by drawing two random numbers from a normal distribution of zero mean and σ standard deviation. We then calculated the grand-canonical weight $p = \exp(-[\Delta E - \mu\Delta N]/k_B T)$, where μ is the chemical potential, $\Delta E, \Delta N$ are the changes in energy and number of particles in the system by the proposed update of the fields, and k_B is the Boltzmann constant. In accordance with the Metropolis–Hastings algorithm, the proposed field update was accepted if $p > 1$ or if




$p \leq x$, where x was a random number drawn from a $[0, 1]$ uniform distribution. Full details can be found in appendix C.

2.2. Data processing

To prepare the output of the Monte Carlo simulations for training we applied the following steps, laid out in figure 1, to each field. First, we simulated the short time-of-flight (TOF) expansion, as performed in the experiment, for 5.3 ms which increased the size of the vortex cores. This involved numerically solving the time-dependent Gross–Pitaevskii equation (GPE) [36] using split-step Fourier methods [37]. We performed simulations in 2D and ignored the interactions as they become negligible after the rapid expansion of the cloud in the z -direction which occurs in the experiment. Simulating in 2D captured the important physics while being significantly less computationally demanding than 3D simulations; see appendix D for further discussion.

Next, we blurred the field using a Gaussian filter of spatial width (standard deviation) σ_1 and constructed a composite field in the density-phase representation using the absolute value of the propagated (unblurred) field and the phase of the blurred field. This coarse-graining [38] removed closely-bound vortex pairs and their associated density dips (length scales smaller than σ_1) which allowed us to detect the unbound vortices that are responsible for the superfluid transition, even as bound vortex-antivortex pairs proliferate. From this composite field, we identified vortices by checking for the characteristic $\pm 2\pi$ phase-winding at each 2×2 plaquette of sites which yielded a binary vortex map of shape 280×280 . We did not distinguish between vortices and antivortices since it is impossible to tell the direction of circulation solely from a density image. Then, we downsampled the vortex map by a factor of four to match the pixel size of our experimental images and applied a mask to restrict vortices to within the central region of the density profile to capture the bulk properties of the system. To emulate the finite resolution of the experimental imaging system, we applied a second Gaussian filter of spatial width σ_2 to the composite field and calculated the squared norm to obtain the density. We applied the same factor-of-four downsampling as with the vortex map, but used the mean values instead of the maximum values of the corresponding 4×4 plaquettes in the original image. Then, we normalized density to the range $[0, 1]$ and added Gaussian white noise of standard deviation $\sigma = 0.03$ to ensure our model was robust against the typical level of background noise present in experimental images. Figure 2 shows two of the simulated images after undergoing processing in comparison with two experimental images.



3. CNN design

We designed our CNN with the pixel size of the experimental images (around $1 \mu\text{m}$) in mind. Our CNN inputs a 70×70 grayscale image of density, with pixel size $1 \mu\text{m}^2$, and passes it through four successive convolutional layers, outputting a vortex map, denoted \tilde{Y} , which is a tensor of the same size where the value of each cell is the probability of a vortex being present in the equivalent cell of the input (see figure 3). The ground truth, denoted Y , is also of the same size, but is restricted to the values 0 and 1. The number of convolutional layers and filters was iteratively optimized to capture vortex patterns with high accuracy while mitigating overfitting to the training data. This refinement ensured robust generalization to the experimental data while minimizing the computational overhead associated with over-parameterization. Notably, omitting fully connected layers significantly reduced memory requirements, improving the model's efficiency and adaptability for deployment.

We used a weighted binary cross-entropy loss function which is commonly used in binary classification:

$$\mathcal{L}(\tilde{Y}, Y) = \sum_{\text{batch}} \sum_{i,j} \left[-w_1 Y_{ij} \log \tilde{Y}_{ij} - (1 - Y_{ij}) \log (1 - \tilde{Y}_{ij}) \right], \quad (2)$$

where the sum i, j is over the grid positions. Since our training set contained significant class imbalance (on average, only 0.32% of sites had vortices), we introduced the w_1 hyperparameter to encourage more positive predictions. Without this, the model often fell into local minima of the loss function corresponding to predicting no vortices at all. Further details of the CNN architecture and training choices can be found in appendix E.

4. Results

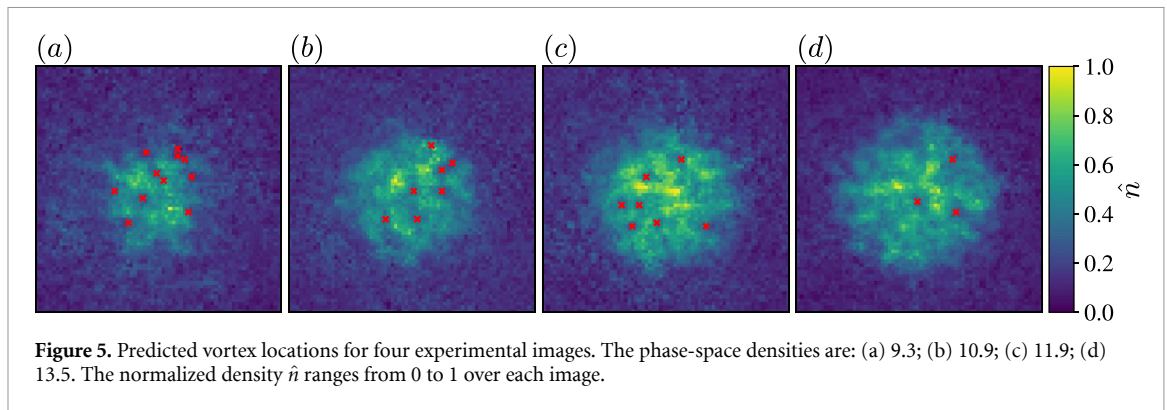
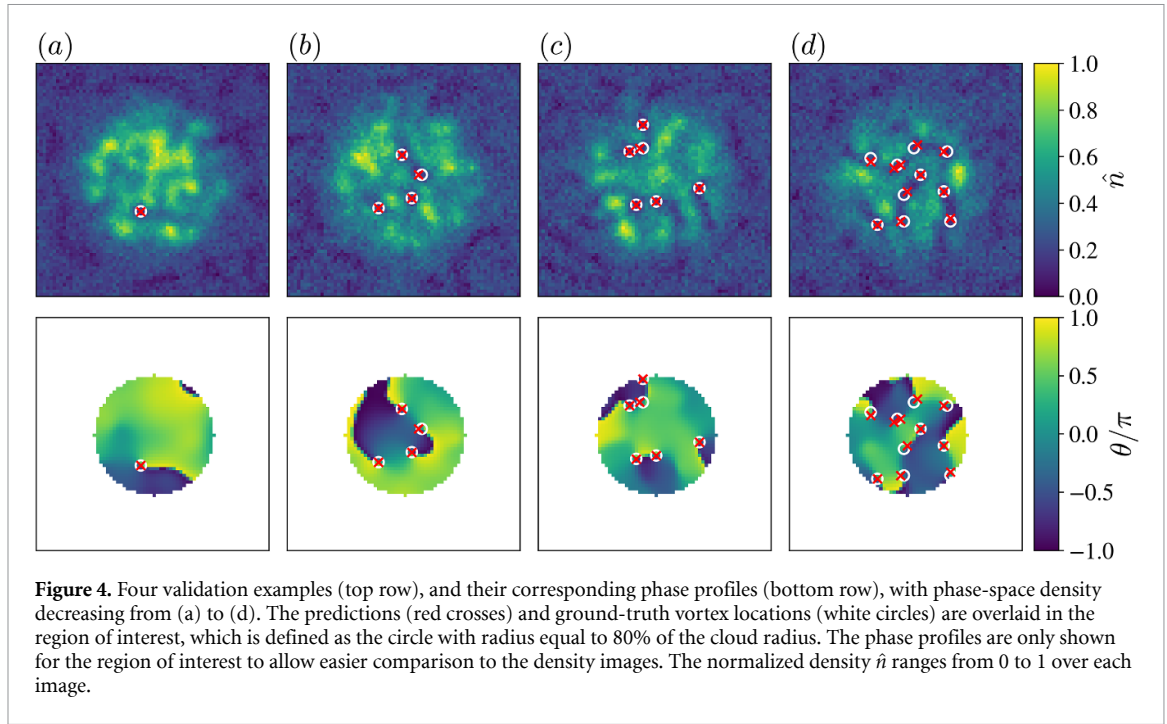
4.1. Validation

Our validation set consisted of 1000 images (5% of the training data) and their accompanying vortex locations. To evaluate our model's predictions on these, we used precision, recall and F1 score which are standard metrics in binary classification and allowed for off-by-one errors where the prediction and ground-truth vortex locations were in adjacent pixels. Metric definitions and justifications for accepting off-by-one errors can be found in appendix F.

We found the precision, recall and F1 scores to be 0.81, 0.65 and 0.72, respectively. We deliberately configured our model to be precision-dominant so we could be confident about any positive predictions it makes on experimental images. Examining figure 4 justifies our leniency with off-by-one errors, and indicates that the model is highly proficient at locating vortices in out-of-sample simulated data. Assessing the phase profiles, we see that the model can distinguish density dips caused by vortices (phase singularities) from those caused by phonons. We chose not to validate the model against ordered noise patterns such as stripes or rings since efficient algorithms to remove such noise patterns are readily available [39].

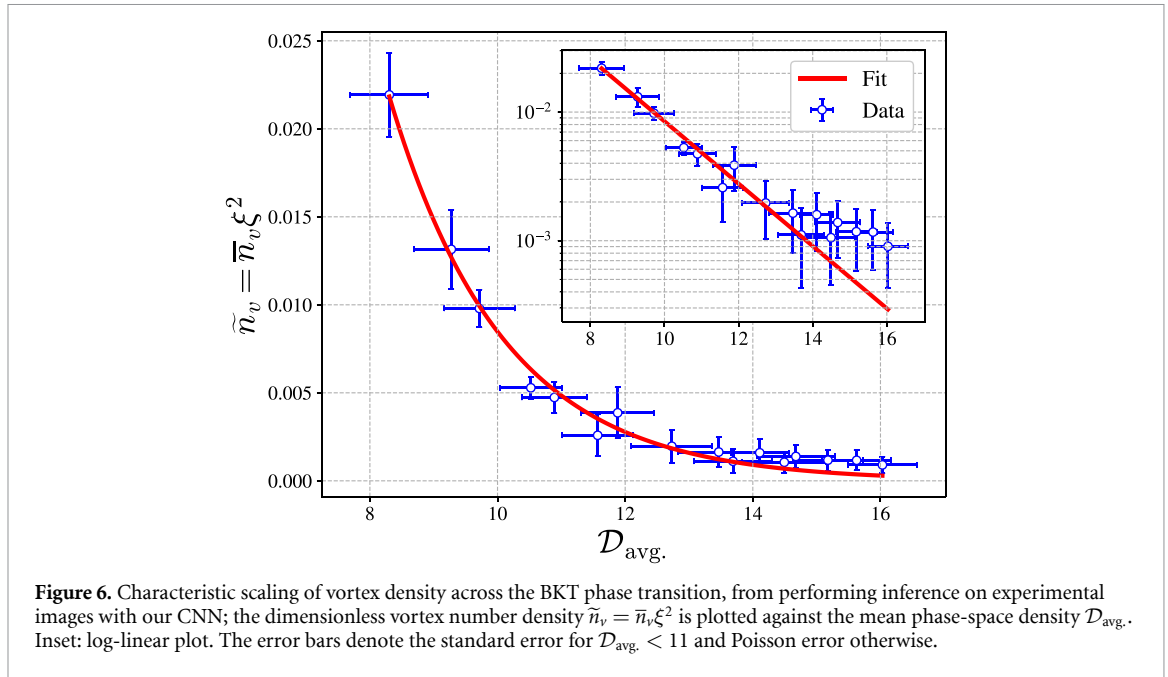
4.2. Predictions on experimental images

The experimental dataset had 876 density images, with mean phase-space densities ranging from 8.3 to 16.0, and required preprocessing before we could make vortex predictions. First, we removed the random



variation of the position of the atom cloud by recentering the images on their center of mass. Then, at each mean phase-space density, we estimated the quasi-condensate radius r_{QC} by fitting the radial density profile, and defined the region of interest as a circle of radius $0.8r_{QC}$ as in [25]. This avoided edge effects where the physics is fundamentally different. We found that reasonable predictions required keeping the experimental images unnormalized as a result of the low experimental resolution and two fundamental differences in the appearances of vortices that were not captured by the simulation. Firstly, as phase-space density drops, the fraction of atoms in the non-condensed phase grows which reduces the contrast of density dips. Secondly, at lower phase-space densities, the reduced detected atom number leads to a lower signal-to-noise ratio. As such, the signal emanating from potential vortex locations must be amplified in order for the method to work. The combination of using unnormalized experimental data and applying image contrast enhancement with morphological operators (see appendix D) achieved this. The model's predictions are shown in figure 5.

We could not verify the accuracy of detecting any particular vortex since their positions are not known *a priori* in the experimental images. However, we analyzed how the dimensionless vortex number density $\tilde{n}_v = \bar{n}_v \xi^2$, where \bar{n}_v is the mean vortex number density and ξ is the healing length, scales with mean phase-space density \mathcal{D}_{avg} . The use of \tilde{n}_v accounts for the variation of vortex size with \mathcal{D}_{avg} , as ξ^2 represents the characteristic area associated with a vortex. The functional form is expected to follow $\tilde{n}_v \propto \exp(-b\mathcal{D}_{avg})$ with exponential coefficient $b = 0.56(5)$ obtained via matter-wave interferometry [25] and supported by simulations. Our fit, shown in figure 6, yielded $b_{exp} = 0.56(2)$ which is consistent with results using the matter-wave interferometric approach. The presence of the rapid vortex number increase gives us confidence that we have detected vortices, not phonons, since phonon amplitude grows slowly with decreasing \mathcal{D}_{avg} and cannot alone explain the phase transition from a superfluid to a normal fluid.



5. Summary & outlook

We have developed a deep learning-based vortex detection method which can infer vortex locations in density images of 2D Bose gases. The detector is a CNN for which the input is a density image, and the output is a corresponding map of vortex locations. We used Monte Carlo simulations to obtain a synthetic dataset on which to train, circumventing the time-consuming and error-prone task of labeling vortices by hand, and fulfilling the requirement of a large training set for performant machine learning. We have shown that our model performs well when evaluated against out-of-sample simulated images scoring 0.81, 0.65 and 0.72 in precision, recall and F1 metrics, respectively. Further, we have been able to generalize the CNN’s predictive power to experimental images where vortices cannot be identified by simple algorithms or the human eye. These predictions were validated against a standard matter-wave interferometric benchmark from previous work, successfully reproducing the ground-truth vortex number scaling. Our results demonstrate that the CNN reliably distinguishes vortices from phonons, marking a significant step toward on-demand, system-wide vortex detection in atomic 2D Bose gases.

With improved image quality, the vortices will be more readily detectable, allowing simple and efficient models—such as the one presented in this paper—to make predictions with even greater confidence and accuracy. Furthermore, transformers can be used to enhance existing CNN-based approaches by integrating global context, better capturing the relationships between vortices, though this comes with increased computational demands. Looking ahead, applications to real-time tracking of free vortices as described in [40–42], or vortex pairing dynamics as in [38] will become more feasible. We hope such improvements will allow this approach to extend to bilayer systems such as those in [23, 25, 27], where inter-layer interactions are expected to give rise to correlated free vortices in the two layers, which can be distinguished from other vortices by density depletions of different magnitude.

Finally, we anticipate that simulation-supported, deep learning-based vortex detection algorithms, like the one presented in this paper, will be a valuable tool not only in the study of quantum gases, but in programmable quantum simulators in general, a quantum technology experiencing rapid growth in recent years [43, 44]. These algorithms are fast and efficient, and since the underlying simulation can be adjusted to match different scenarios, they are also very adaptable.

Data availability statement

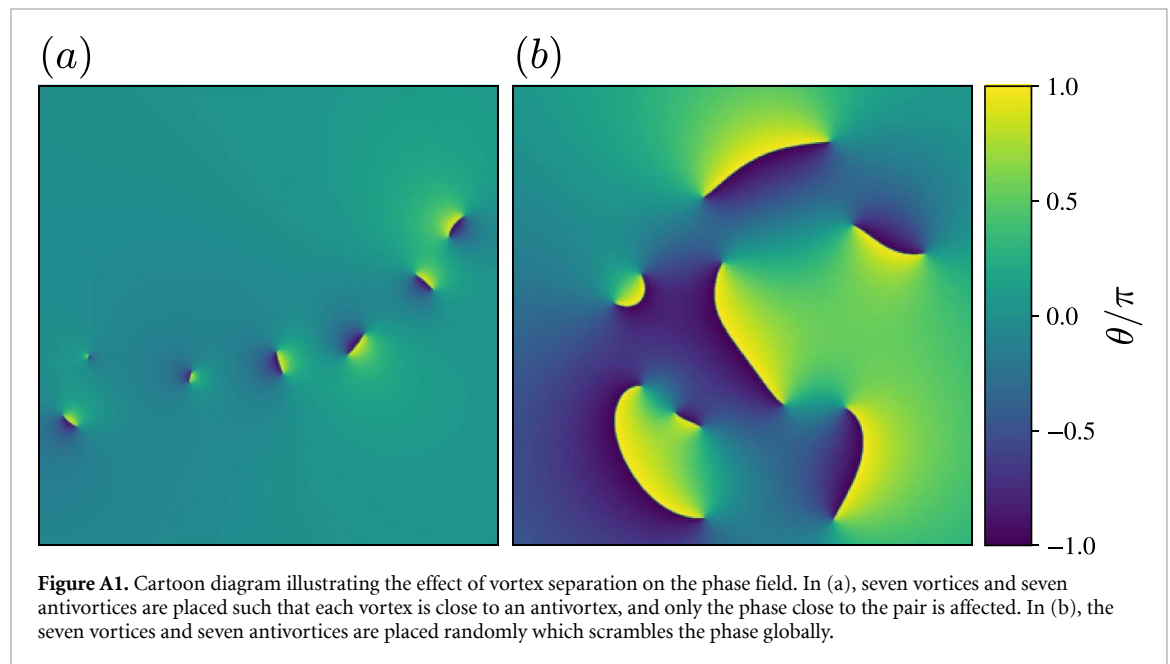
Example preprocessing, training, and inference notebooks are available at [45] alongside a trained model. The data cannot be made publicly available upon publication because no suitable repository exists for hosting data in this field of study. The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgments

MS received support from the John O'Connor Research Fund, St Peter's College, Oxford. The data analyzed was from experimental work supported by EPSRC Grant EP/X024601/1. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used in this research.

Appendix A. The BKT transition in 2D Bose gases

In uniform 2D Bose gases, fluctuations (phonons) destroy true long-range order at any finite temperature, thus inhibiting the formation of a Bose–Einstein condensate, but the system still exhibits superfluid behavior at sufficiently low temperatures. The phase transition between the superfluid phase at low temperature and the normal phase at high temperature was explained by Berezinskii [46], Kosterlitz, and Thouless [47], who identified the key microscopic mechanism involving vortices which drives the transition. The transition is between a superfluid, quasi-ordered phase where correlations decay algebraically to a disordered, normal fluid phase where correlations decay exponentially. At low temperatures, unbound vortices are strongly suppressed due to their large energy cost compared to the thermal energy available, instead closely-bound vortex-antivortex pairs form [48]. The presence of these closely-bound pairs still allows superfluid behavior as the circulation of the pairs cancels for contours which enclose the pair. Such vortex pairs only affect the phase significantly on length scales comparable with the pair separation, therefore the local phase of the quantum system remains well-correlated over larger length scales. Upon further temperature increase, the vortices in these pairs gain the energy to move further apart until the system reaches the critical temperature at which they unbind, resulting in free vortices moving in the system. This unbinding scrambles the phase over all relevant length scales, i.e. the quasi-long-range order makes way for disorder, corresponding to a transition from the superfluid phase into the normal phase. Figure A1 illustrates the effect of vortex separation on the phase-field.



Appendix B. Experimental details

We prepare 2D degenerate Bose gases of ^{87}Rb atoms in a cylindrically-symmetric trap having strong confinement along the z -direction. In this work, we created vertical confinement using the radio-frequency dressing techniques described in [25, 49, 50], with a vertical trap frequency of $\omega_z/2\pi = 1.2$ kHz. This resulted in the dimensionless 2D interaction strength $\tilde{g} = \sqrt{8\pi}a_s/\ell_0 = 0.08$, where a_s is the 3D s-wave scattering length, $\ell_0 = \sqrt{\hbar/(m\omega_z)}$ is the harmonic oscillator length along z for an atom of mass m , and \hbar is the reduced Planck constant. We imposed additional optical trapping in the horizontal plane using a ring-shaped, strong, off-resonant laser beam to realize near-homogeneous 2D systems [23]. We loaded N

atoms into the trap at a temperature $T = 50$ nK, set by forced evaporation. We varied N to cover a broad range of the phase-space density $\mathcal{D} = nh^2/(2\pi mk_B T)$, where n is the average 2D density. For all parameters employed, the quasi-2D conditions $\hbar\omega_z > k_B T$ and $\hbar\omega_z > \mu$ were satisfied.

Appendix C. Simulation details & verification

To ensure thermal equilibrium and uncorrelated fields, we initially ran the algorithm for 200 000 updates per site, and then began sampling fields with a further 5000 updates per site occurring between each field sample. We tuned σ such that the mean acceptance rate of field updates was near 0.25, which yields optimal convergence [51, 52]. We collected 2500 fields at eight phase-space densities evenly spaced in the range $\mathcal{D} \in [8.0, 11.2]$ by keeping the temperature constant at $T = 50$ nK and varying μ . We expected the phase transition to occur at $\mathcal{D} \approx 11$ for our finite-size system which we verify by direct analysis of the two-point phase correlation function in figure C1.

We can see from figure C1(a) that the BKT transition occurs around $\mathcal{D} \approx 11$ as evidenced by the narrow crossover region between the root-mean-square errors of algebraic and exponential fits to the correlation function. In figure C1(b), we see that the crossover region places the algebraic exponent η , defined when fitting $r^{-\eta}$ to the correlation function, just below the universal value of 0.25 found in the thermodynamic limit—a typical result for finite-sized systems [25].

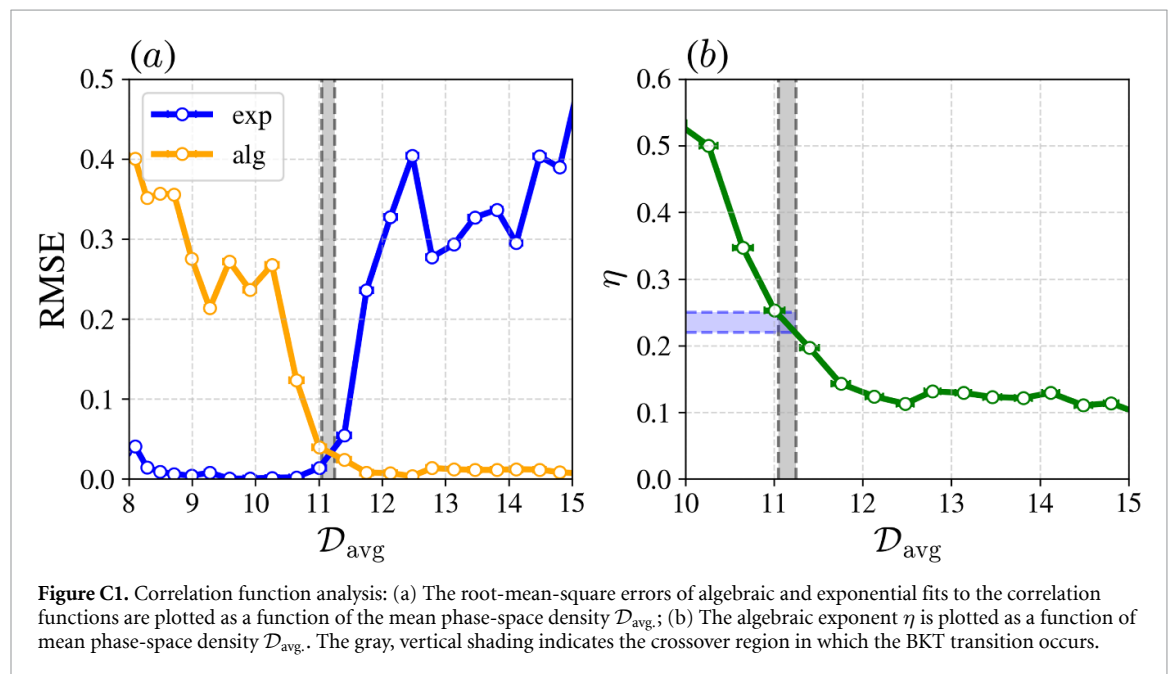
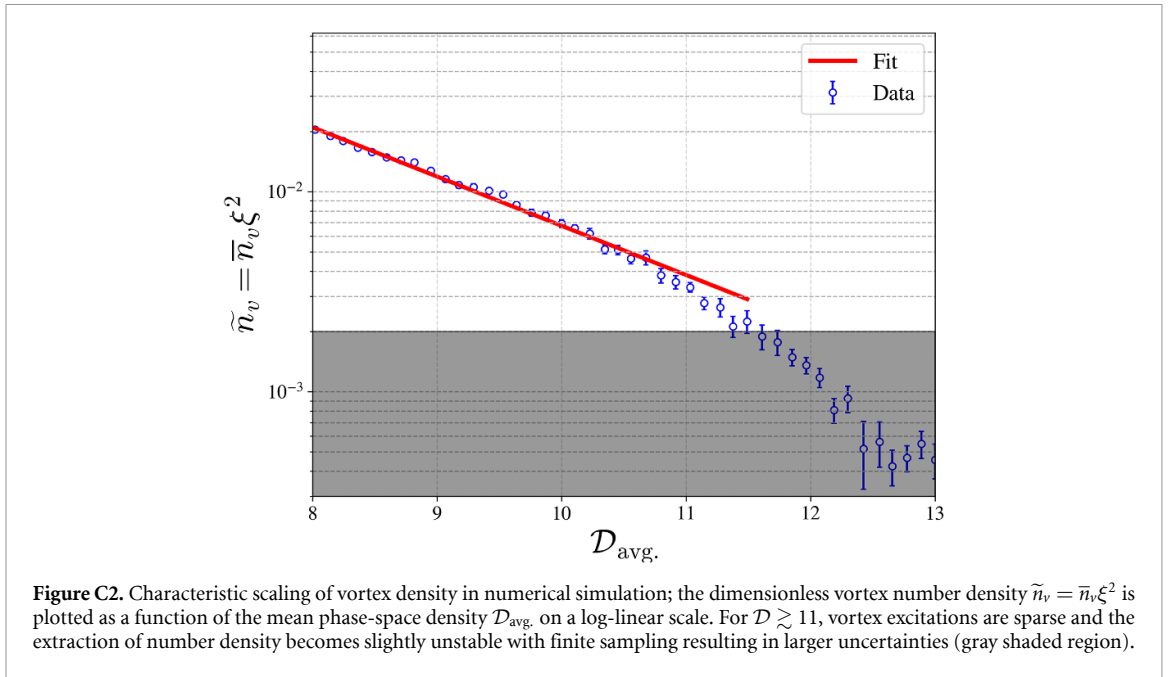


Figure C1. Correlation function analysis: (a) The root-mean-square errors of algebraic and exponential fits to the correlation functions are plotted as a function of the mean phase-space density \mathcal{D}_{avg} ; (b) The algebraic exponent η is plotted as a function of mean phase-space density \mathcal{D}_{avg} . The gray, vertical shading indicates the crossover region in which the BKT transition occurs.

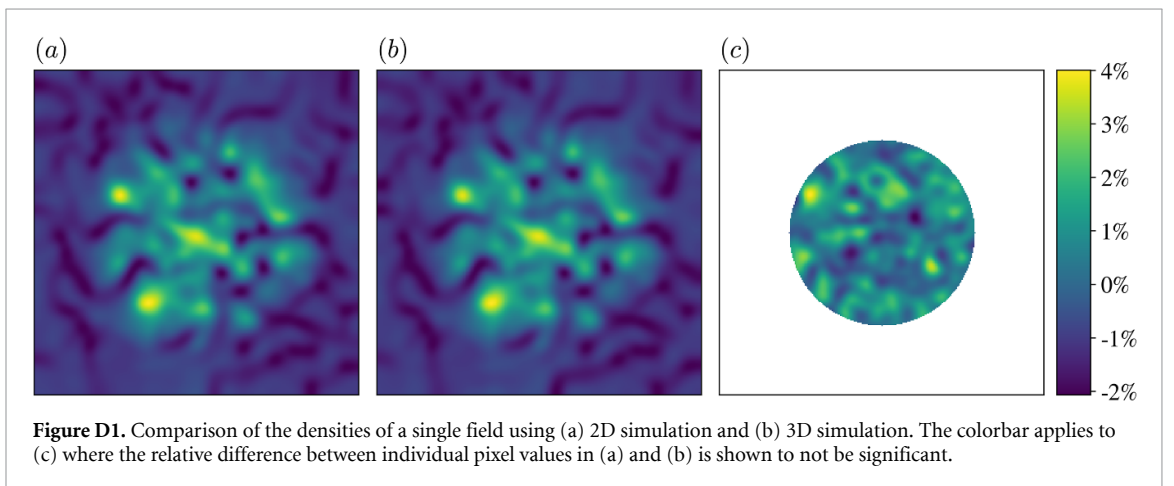
We verified that our simulation reproduced the scaling of dimensionless vortex number density \tilde{n}_v , with mean phase-space density \mathcal{D}_{avg} , previously observed in [25]. Figure C2 shows the result on a log-linear scale, with our fit region restricted to the range of mean phase-space densities present in our training data ($\mathcal{D}_{\text{avg}} \in [8.0, 11.2]$). We found the exponential fit coefficient $b_{\text{sim.}} = 0.57(1)$ to be consistent with previous experimental results. We ignored the region $\tilde{n}_v < 2 \times 10^{-3}$ as the vortex counts are extremely low, averaging less than a single vortex per image.

To ensure the lattice accurately represented the continuum limit, we compared the site separation $\delta l = 0.25 \mu\text{m}$ to two characteristic length scales for a Bose gas: (1) the healing length $\xi = \hbar/\sqrt{2mg\tilde{n}}$ which is the characteristic size of the vortex core, and (2) the thermal de Broglie wavelength $\lambda = h/\sqrt{2\pi k_B T}$. We found, empirically, that our healing length falls in the range $\xi \in [0.59, 0.99] \mu\text{m}$, and at $T = 50$ nK the thermal de Broglie wavelength is $\lambda = 0.85 \mu\text{m}$. Thus, $\delta l \lesssim \xi, \lambda$, and our lattice approximation is justified.

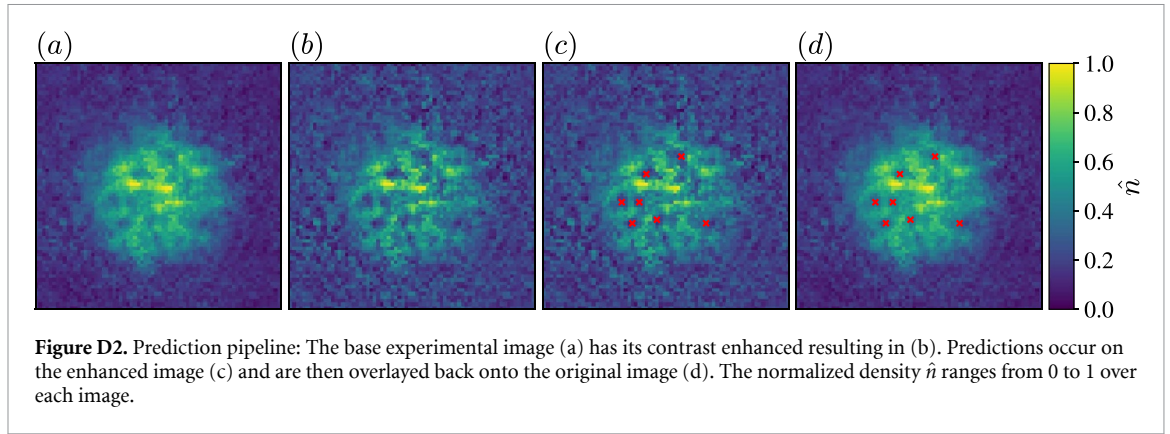


Appendix D. Data processing methods

We justified the choice to simulate the time-of-flight expansion in 2D, as opposed to in 3D, by plotting the density images that arose in either case. For the 3D simulation, the density was averaged along z in a region of width $10 \mu\text{m}$ as in the experiment. We see, in figure D1, that the density images are qualitatively indistinguishable, and by subtracting one from the other, we see that the relative difference of any given pixel is at most a few percentage points.



We applied contrast enhancement, inspired by [53], to the experimental images to aid prediction. In brief, the higher density (bright) and lower density (dark) areas were identified using the morphological operators top-hat and black-hat, respectively. The bright areas were added to, and the dark areas were subtracted from, the original image to form the enhanced image. The effect of this enhancement is shown in figure D2.



Appendix E. Neural network architecture and training details

Our CNN has four successive convolutional layers of which the first three use the rectified linear unit (ReLU) activation function and the final layer uses the sigmoid activation function to obtain a probability between zero and one. Each convolutional layer, barring the final one, is followed by a batch normalization layer to stabilize and speed up training [54]. The filter sizes and number of filters in each layer are shown in table E1. Including the batch normalization layers, there are 24 406 learnable parameters.

Before training, we normalized our combined training and validation dataset to have pixel values in the range [0, 1]. Then, using the loss function hyperparameter $w_1 = 3.3$, a learning rate of $\alpha = 0.0001$, a batch size of 32 and the Adam optimization algorithm [55], we trained for 10 epochs. The training PC was equipped with 64 GB of RAM and a single NVIDIA Titan Xp GPU, resulting in a training time of around 90 s, and inference of 1000 images in 350 ms.

Table E1. Neural Network Architecture. Batch normalization layers which occur after every convolutional layer are not shown but contain additional parameters.

Layer	No. Filters	Filter Size	Stride	No. Parameters
Conv1	30	5×5	1	780
Conv2	25	5×5	1	18 775
Conv3	20	3×3	1	4520
Conv4	1	3×3	1	181

Appendix F. Validation metrics

Our evaluation metrics were based on precision and recall which are standard in binary classification. Denoting the number of true positives as TP, false negatives as FN, and similarly for TN and FP, we define precision and recall as

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (\text{F1})$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}. \quad (\text{F2})$$

The F1 score is defined as the harmonic mean of precision and recall:

$$\text{F1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (\text{F3})$$

We allowed for off-by-one errors in the above metrics, and justified this choice as follows. The characteristic size of an *in-situ* vortex is given by its healing length which we found to be in the range $\xi \in [0.59, 0.99] \mu\text{m}$. After TOF expansion, the size of the vortex core increases by a factor of two, verified by GPE simulations. Given our pixel size ($1 \mu\text{m}$) was of the same order of magnitude, we expected individual

vortices to have physical extent over multiple pixels. Additionally, the ground-truth locations contained their own ambiguity due to the factor-of-four downsampling. As such, we expected the model to incur off-by-one errors, where the prediction and ground-truth vortex locations are in adjacent pixels, and for the model to sometimes predict the same vortex in adjacent cells. Since we were only interested in isolated vortices, the off-by-one errors were safely ignored, and the adjacent predictions were fixed by post-processing.

The post-processing was a custom non-max suppression (NMS) algorithm. NMS is a common post-processing stage in the field of object detection and is typically used to choose a single bounding box for an object when many are generated [56, 57]. Here our NMS algorithm kept only the highest probability vortex location in any given pair of predictions on adjacent cells when both prediction probabilities were greater than 1/2.

Appendix G. Density profiles of detected vortices

To further verify that the CNN had only detected vortices in the experimental images, we calculated the width of the density dips at the predicted vortex locations. We averaged the 7×7 regions where vortices were predicted for a given phase-space density (see figure G1(a)), and fitted a Gaussian to the corresponding radial density profile to estimate the vortex size r_v , defined as the half-width at half-maximum of the fit (see figure G1(b)). We found $r_v = 1.84(3) \mu\text{m}$, which is close to the value obtained from zero-temperature GPE simulations of TOF expansion dynamics of a single vortex in a system of comparable density.

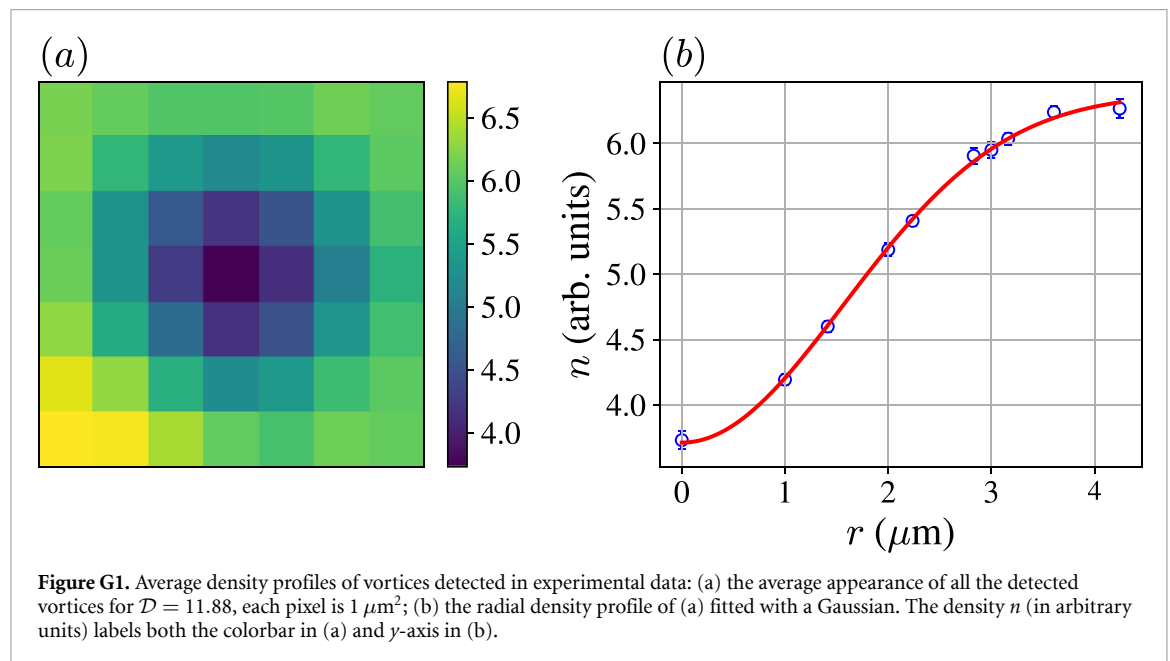


Figure G1. Average density profiles of vortices detected in experimental data: (a) the average appearance of all the detected vortices for $\mathcal{D} = 11.88$, each pixel is $1 \mu\text{m}^2$; (b) the radial density profile of (a) fitted with a Gaussian. The density n (in arbitrary units) labels both the colorbar in (a) and y -axis in (b).

ORCID iDs

M Sesodia  <https://orcid.org/0009-0000-5217-6639>

S Sunami  <https://orcid.org/0000-0002-0969-9909>

E Chang  <https://orcid.org/0009-0000-8794-5581>

E Rydow  <https://orcid.org/0000-0001-5140-9461>

C J Foot  <https://orcid.org/0000-0001-8822-919X>

A Beregi  <https://orcid.org/0000-0001-5647-4051>

References

- [1] Gessner O and Vilesov A F 2019 *Annu. Rev. Phys. Chem.* **70** 173–98
- [2] Tang Y, Guo W, Kobayashi H, Yui S, Tsubota M and Kanai T 2023 *Nat. Commun.* **14** 2941
- [3] Roditchev D et al 2015 *Nat. Phys.* **11** 332–7
- [4] Llorens J B et al 2020 *Phys. Rev. Res.* **2** 013329
- [5] Genevet P, Barland S, Giudici M and Tredicce J R 2010 *Phys. Rev. Lett.* **104** 223902
- [6] Situ G and Fleischer J W 2020 *Nat. Photon.* **14** 517–22

- [7] Panico R et al 2021 *Phys. Rev. Lett.* **127** 047401
- [8] Choi D et al 2022 *Phys. Rev. B* **105** L060502
- [9] Matthews M R, Anderson B P, Haljan P C, Hall D S, Wieman C E and Cornell E A 1999 *Phys. Rev. Lett.* **83** 2498–501
- [10] Hadzibabic Z, Krüger P, Cheneau M, Battelier B and Dalibard J 2006 *Nature* **441** 1118–21
- [11] Zwierlein M W, Abo-Shaer J R, Schirotzek A, Schunck C H and Ketterle W 2005 *Nature* **435** 1047–51
- [12] Liu X P, Yao X C, Deng Y, Wang Y X, Wang X Q, Li X, Chen Q, Chen Y A and Pan J W 2021 *Phys. Rev. Res.* **3** 043115
- [13] Lee K, Kim S, Kim T and Shin Y 2024 *Nat. Phys.* **20** 1570–4
- [14] Williams R A, Al-Assam S and Foot C J 2010 *Phys. Rev. Lett.* **104** 050404
- [15] Sharma R, Rey D, Longchambon L, Perrin A, Perrin H and Dubessy R 2024 *Phys. Rev. Lett.* **133** 143401
- [16] Chomaz L, Corman L, Bienaimé T, Desbuquois R, Weitenberg C, Nascimbène S, Beugnon J and Dalibard J 2015 *Nat. Commun.* **6** 6162
- [17] Rabga T, Lee Y, Bae D, Kim M and Shin Y 2023 *Phys. Rev. A* **108** 023315
- [18] Posazhennikova A 2006 *Rev. Mod. Phys.* **78** 1111–34
- [19] Maccari I, Defenu N, Benfatto L, Castellani C and Enss T 2020 *Phys. Rev. B* **102** 104505
- [20] Ketterle W, Durfee D S and Stamper-Kurn D 1999 Making, probing and understanding Bose-Einstein condensates *Bose-Einstein Condensation in Atomic Gases* (IOS Press) pp 67–176
- [21] Hueck K, Luick N, Sobirey L, Siegl J, Lompe T, Moritz H, Clark L W and Chin C 2017 *Opt. Express* **25** 8670–9
- [22] Choi J y, Seo S W, Kwon W J and Shin Y I 2012 *Phys. Rev. Lett.* **109** 125301
- [23] Sunami S, Singh V P, Rydow E, Beregi A, Chang E, Mathey L and Foot C J 2024 Detecting phase coherence of 2D Bose gases via noise correlations (arXiv: 2406.03491)
- [24] Stock S, Hadzibabic Z, Battelier B, Cheneau M and Dalibard J 2005 *Phys. Rev. Lett.* **95** 190403
- [25] Sunami S, Singh V P, Garrick D, Beregi A, Barker A J, Luksch K, Bentine E, Mathey L and Foot C J 2022 *Phys. Rev. Lett.* **128** 250402
- [26] Sunami S, Singh V P, Garrick D, Beregi A, Barker A J, Luksch K, Bentine E, Mathey L and Foot C J 2023 *Science* **382** 443–7
- [27] Beregi A, Foot C and Sunami S 2024 *AVS Quantum Sci.* **6** 030501
- [28] Metz F, Polo J, Weber N and Busch T 2021 *Mach. Learn.: Sci. Technol.* **2** 035019
- [29] Kim M, Kwon J, Rabga T and Shin Y 2023 *Mach. Learn.: Sci. Technol.* **4** 045017
- [30] Prokof'ev N, Ruebenacker O and Svistunov B 2001 *Phys. Rev. Lett.* **87** 270402
- [31] Blakie P B, Bradley A S, Davis M J, Ballagh R J and Gardiner C W 2008 *Adv. Phys.* **57** 363–455
- [32] Kiehn H, Singh V P and Mathey L 2022 *Phys. Rev. A* **105** 043317
- [33] Singh V P and Mathey L 2022 *New J. Phys.* **24** 073024
- [34] Metropolis N, Rosenbluth A W, Rosenbluth M N, Teller A H and Teller E 1953 *J. Chem. Phys.* **21** 1087–92
- [35] Hastings W K 1970 *Biometrika* **57** 97–109
- [36] Dalfovo F, Giorgini S, Pitaevskii L P and Stringari S 1999 *Rev. Mod. Phys.* **71** 463–512
- [37] Javanainen J and Ruostekoski J 2006 *J. Phys. A: Math. Gen.* **39** L179–84
- [38] Foster C J, Blakie P B and Davis M J 2010 *Phys. Rev. A* **81** 023623
- [39] Song B, He C, Ren Z, Zhao E, Lee J and Jo G B 2020 *Phys. Rev. Appl.* **14** 034006
- [40] Andersen M F, Ryu C, Cladé P, Natarajan V, Vaziri A, Helmerson K and Phillips W D 2006 *Phys. Rev. Lett.* **97** 170406
- [41] Navarro R, Carretero-González R, Torres P J, Kevrekidis P G, Frantzeskakis D J, Ray M W, Altuntaş E and Hall D S 2013 *Phys. Rev. Lett.* **110** 225301
- [42] Koukouloyannis V, Voyatzis G and Kevrekidis P G 2014 *Phys. Rev. E* **89** 042905
- [43] Schäfer F, Fukuhara T, Sugawa S, Takasu Y and Takahashi Y 2020 *Nat. Rev. Phys.* **2** 411–25
- [44] Ebadi S et al 2021 *Nature* **595** 227–32
- [45] Sesodia M 2024 CNN-Based vortex detection in atomic 2D Bose gases in the presence of a phononic background (available at: https://github.com/magnus-sesodia/2d_vortex_detection)
- [46] Berezinskiĭ V 1972 *Sov. Phys. JETP* **34** 610
- [47] Kosterlitz J M and Thouless D J 1973 *J. Phys. C: Solid State Phys.* **6** 1181
- [48] Kosterlitz J M 1974 *J. Phys. C: Solid State Phys.* **7** 1046–7
- [49] Barker A J, Sunami S, Garrick D, Beregi A, Luksch K, Bentine E and Foot C J 2020 *J. Phys. B: At. Mol. Opt. Phys.* **53** 155001
- [50] Barker A J, Sunami S, Garrick D, Beregi A, Luksch K, Bentine E and Foot C J 2020 *New J. Phys.* **22** 103040
- [51] Gelman A, Gilks W R and Roberts G O 1997 *Ann. Appl. Probab.* **7** 110–20
- [52] Roberts G O and Rosenthal J S 2001 *Stat. Sci.* **16** 351–67
- [53] Widyantara I M O, Asana I M D P, Wirastuti N and Adnyana I B P 2016 Image enhancement using morphological contrast enhancement for video based image analysis 2016 *Int. Conf. Data Softw. Eng. (ICoDSE)*
- [54] Ioffe S 2015 Batch normalization: accelerating deep network training by reducing internal covariate shift (arXiv: 1502.03167)
- [55] Kingma D P and Ba J 2017 Adam: a method for Stochastic Optimization (arXiv: 1412.6980)
- [56] Girshick R, Donahue J, Darrell T and Malik J 2014 Rich feature hierarchies for accurate object detection and semantic segmentation (arXiv: 1311.2524)
- [57] Redmon J, Divvala S, Girshick R and Farhadi A 2016 You only look once: unified, real-time object detection (arXiv: 1506.02640)