

Deep Denoising for Scientific Discovery: A Case Study in Electron Microscopy

Sreyas Mohan^{*}, Ramon Manzorro[†], Joshua L. Vincent[‡], Binh Tang[‡], Dev Y. Sheth[§], Eero P. Simoncelli[¶], David S. Matteson[§], Peter A. Crozier[‡], and Carlos Fernandez-Granda^{||}

Abstract. Denoising is a fundamental challenge in scientific imaging. Deep convolutional neural networks (CNNs) provide the current state of the art in denoising natural images, where they produce impressive results. However, their potential has been inadequately explored in the context of scientific imaging. Denoising CNNs are typically trained on real natural images artificially corrupted with simulated noise. In contrast, in scientific applications, noiseless ground-truth images are usually not available. To address this issue, we propose a simulation-based denoising (SBD) framework, in which CNNs are trained on simulated images. We test the framework on data obtained from transmission electron microscopy (TEM), an imaging technique with widespread applications in material science, biology, and medicine. SBD outperforms existing techniques by a wide margin on a simulated benchmark dataset, as well as on real data. We analyze the generalization capability of SBD, demonstrating that the trained networks are robust to variations of imaging parameters and of the underlying signal structure. Our results reveal that state-of-the-art architectures for denoising photographic images may not be well adapted to scientific-imaging data. For instance, substantially increasing their field-of-view dramatically improves their performance on TEM images acquired at low signal-to-noise ratios. We also demonstrate that standard performance metrics for photographs (such as PSNR and SSIM) may fail to produce scientifically meaningful evaluation. We propose several metrics to remedy this issue for the case of atomic resolution electron microscope images. In addition, we propose a technique, based on likelihood computations, to visualize the agreement between the structure of the denoised images and the observed data. Finally, we release a publicly available benchmark dataset of TEM images, containing 18,000 examples.

1. Introduction. Imaging technology is an essential tool in many scientific domains. Electron microscopy enables the visualization of atomic structures [79], fluorescence microscopy makes it possible to study cellular processes [37], and telescopes reveal galaxies and other astronomical objects that are light years away [45]. In all these modalities, images are corrupted by noise associated with stochastic processes occurring during signal generation and detection, degrading the information content of the image data. The general goal of denoising is to estimate and restore the information missing from these noisy observations, thus facilitating the extraction of useful scientific information.

In the past decade, convolutional neural networks (CNNs) [34] have achieved state-of-the-art performance in image denoising [75, 13]. However, the potential of this methodology has barely been explored in the context of scientific imaging. In the vast majority of the existing work, noisy data are generated by adding Gaussian noise to clean photographs. The CNNs are then trained to approximate the ground-truth images from these measurements, usually by minimizing mean squared error [75]. Unfortunately, *this paradigm is not adequate for most scientific domains*, where large, labeled datasets of ground-truth clean data are typically not available. To address this issue, we propose a simulation-based denoising (SBD) framework, in which CNNs are trained on simulated images. We

^{*}Center for Data Science, New York University

[†]School for Engineering of Matter, Transport, and Energy, ASU

[‡]Department of Statistics and Data Science, Cornell University

[§]Indian Institute of Technology Madras, India

[¶]Howard Hughes Medical Institute, Center for Neural Science, and Courant Institute of Mathematical Sciences, NYU

^{||}Courant Institute of Mathematical Sciences and Center for Data Science, NYU

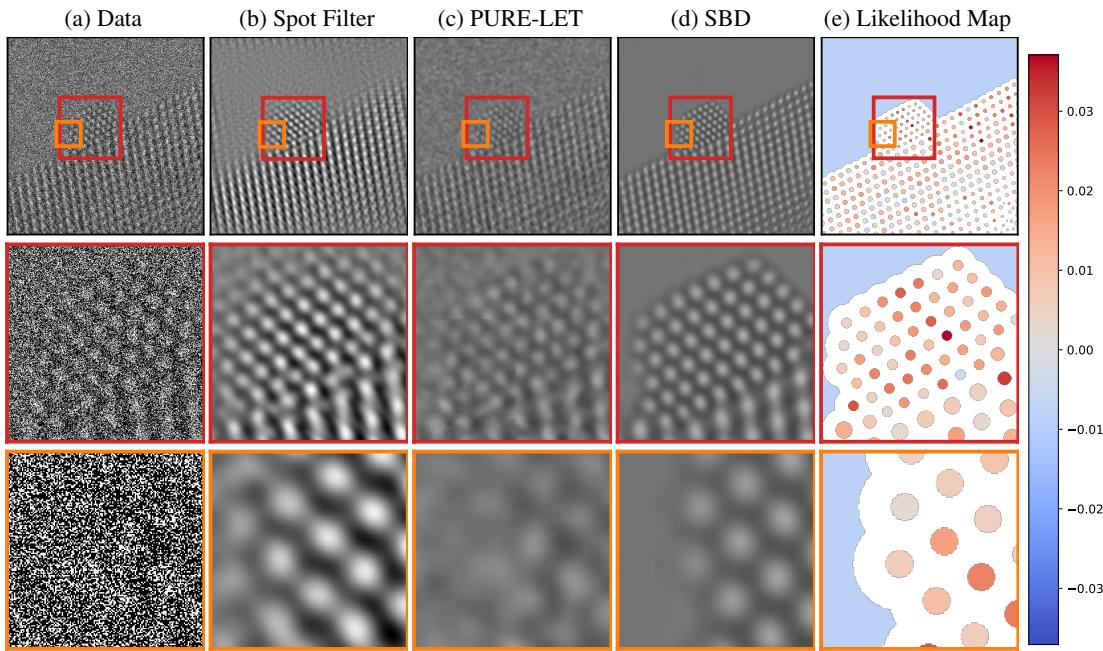


Figure 1.1: Denoising results for real data. (a) An experimentally-acquired atomic-resolution transmission electron microscope image of a CeO₂-supported Pt nanoparticle. A description of the experimental data acquisition is given in Section 4. The average image intensity is 0.45 electrons/pixel (i.e., a large fraction of pixels register zero electrons), which results in an extremely low signal-to-noise ratio. (b) Denoised image obtained via Fourier-based filtering by a domain expert. (c) Denoised image obtained via the wavelet-based PURE-LET method [41]. (d) Denoised image obtained by the proposed simulation-based denoising (SBD) framework. (e) Likelihood map quantifying to what extent the atomic structure identified from the SBD denoised image is consistent with the data (see Section 3). Regions in red are more likely to correspond to atomic columns in the nanoparticle. Regions in blue are more likely to belong to the vacuum.

validate our methodology through a case study in transmission electron microscopy.

Transmission electron microscopy (TEM) is a powerful and versatile characterization technique used to probe the atomic-level structure and composition of a wide range of materials, such as catalysts or semiconductors [64, 67]. The technique has had a huge impact in structural biology, as recognized with the award of the 2017 Nobel Prize in Chemistry [2]. Recent advancements in direct electron detection systems enable experimentalists to image dynamic events at frame rates in the kilohertz range [18, 17]. Imaging at these time scales is critical to advance our understanding of functional materials. In catalytic systems, for example, the chemical transformation process is accompanied by dynamic, atomic-level structural rearrangements which may occur over a time scale spanning tens of milliseconds [65, 21, 33, 36, 14]. Acquiring image series at such high temporal resolution necessarily produces datasets that are severely degraded by shot noise, rendering traditional imaging processing approaches ineffective. It is typically not feasible to reduce the noise content by increasing the intensity

of the incident electron beam, since the high-energy beam can also damage the material when exposed to high doses. Consequently, there is an acute need for novel denoising technology in this domain.

In order to apply the proposed SBD framework to TEM data, we generate a large simulated dataset of TEM images, containing 18,000 examples, and use it to train CNNs for noise removal. This approach outperforms existing techniques by a wide margin on held-out simulated data, as well as on real TEM measurements (see Figures 1.1, 5.2 and 5.6 and Sections 5.2 and 5.4). We perform a thorough analysis of the generalization capability of our models, demonstrating that the CNNs are robust to variations of imaging parameters and of underlying signal structure. Our results indicate that architectures optimized for natural photographic images may have fundamental shortcomings when applied to domain-specific data. For instance, we show that substantially increasing the field-of-view of denoising CNNs has almost no effect on photographs, but produces a significant boost in performance for TEM images. We also demonstrate that standard performance metrics for photographs, such as peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [71], often fail to produce a scientifically-meaningful evaluation of the denoising results. For example, the presence or absence of a single atomic column often results in a negligible change in these metrics. This is highly problematic, because detecting these columns is one of the main motivations for our case study. To remedy this issue, we propose several scientifically-motivated metrics to evaluate our results (see Section 5.3). In addition, we propose a likelihood-based visualization of the agreement between the observed measurements and structures of interest (such as atomic columns) in the denoised image. This visualization can be used to flag denoising artefacts, which may be mistaken for scientifically-relevant structure (see Figure 3.3). Finally, to encourage further development of deep-learning methodologies for scientific imaging, we have made our benchmark dataset of TEM images publicly available online¹. More details on applying the proposed methodology to TEM data, and domain-specific insights derived from the denoised images are described in our companion paper [70].

2. Related work.

Denoising in scientific imaging. A wide variety of denoising methods have been applied across different scientific imaging modalities, including traditional linear filters [52], nonlinear filters [68, 47, 25], wavelet-based methods [12, 55, 78, 46], and sparsity-based approaches [46, 6]. Deep convolutional networks have been shown to outperform all of these approaches in photographic images [75, 13]. The rapidly growing literature on this methodology focuses almost exclusively on photographic images. We are aware of only a few very recent exceptions. In the medical domain, CNN-based denoising has been applied to low-dose computer tomography [27], positron-emission tomography [20] and scintillation-camera data [48]. Refs. [19, 69, 16] apply CNNs to denoise simulated electron microscopy data, without validating on real data. Ref. [44] trains CNNs to denoise Raman scattering microscopy data, using measurements gathered at a higher signal-to-noise ratio (SNR) as ground-truth images. These results showcase the potential of deep denoising for scientific imaging, but also the challenge of gathering adequate datasets to train the deep networks. In this work, we propose to address this challenge by training denoising CNNs on carefully-designed simulated datasets, and validate our approach on experimental measurements.

Unsupervised denoising. Unsupervised denoising is a promising approach for applications where ground-truth images are not available. Unsupervised methods based on wavelets have achieved

¹<https://sreyas-mohan.github.io/electron-microscopy-denoising/>

performance comparable to their supervised counterparts on photographic images [40, 22, 60]. Noise2Noise [35], a deep-learning approach that requires access to pairs of noisy images corresponding to the same underlying signal, has been applied to cryo-electron microscopy [11]. More recent methods can be trained directly on noisy images [30, 5, 32]. Several recent works apply this approach to fluorescence microscopy data [31, 26, 56, 77]. In the case of our TEM data, standard unsupervised methods do not perform as well as the proposed supervised approach (see Section 5.4.1). This is possibly due to the limited number of training data (see Figure 5.8) and to the low input SNR. The SNR of our TEM data (around 3 dB) is orders of magnitude lower than that reported in typical unsupervised denoising works (e.g. around 27 dB for [77]).

Deep Learning for TEM. Deep CNNs have been applied to other image-processing tasks in TEM beyond denoising, see Ref. [15] for a comprehensive review. Ref. [66] proposes a CNN-based method for TEM image super-resolution, wherein CNNs are trained on pairs of low-resolution and high-resolution images acquired experimentally. Ref. [23] applies CNNs to perform segmentation and systematically studies the influence of the design of the training dataset and network architecture on the generalization capabilities of these models. In this work, we provide a similar analysis for denoising. Refs. [42] and [59] propose a CNN-based method to identify structures of interest in TEM images. They train on carefully designed simulated data and show that the model generalizes to real data. Our work provides further evidence that CNNs trained on simulated data can generalize effectively to real measurements.

3. Methodology.

3.1. Simulation-based denoising. Current state-of-the-art deep-learning techniques for denoising photographic images require a training set of ground-truth images [75]. Typically these clean images are corrupted with additive Gaussian noise, and the CNNs are trained to minimize the mean squared error between the network output and the original images. The main obstacle to leveraging this approach in scientific imaging is the lack of ground-truth data; in many applications there is no such thing as a *clean image*. We address this by using a dataset of simulated images to train the CNNs. We call this framework simulation-based denoising (SBD).

Simulation-based denoising (SBD) consists of three stages: simulation of the training set, training of the CNNs using the simulated data, and inference on the real data (see Figure 2.1 for an overview of the methodology). In order to generate the training set, we simulate clean images $x_1, \dots, x_N \in \mathbb{R}^M$ (where M is the number of pixels) according to a predefined physical model. These clean images are then corrupted using a noise model, which can follow a predefined model or be learned from the data, to generate the simulated noisy data. We provide a detailed account of how we generate the simulated dataset for our case study in Sections 4 and A.1, and of the noise model in Section 4.3. Let $Y(x_i)$ denote the random vector representing the noisy image corresponding to the clean simulated image x_i and let $y(x_i)$ represent a realization of $Y(x_i)$. We parameterize the denoising function as a CNN $f_\theta : \mathbb{R}^M \rightarrow \mathbb{R}^M$ where the parameters θ are the weights of the network. To find a good denoising function f_θ , we minimize a loss function $\mathcal{L} : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}$ which quantifies how close the estimate from the CNN $f_\theta(y(x_i))$ is to the clean image x_i . In our case study, we use mean squared error, which is a standard choice in CNN-based denoising [75]. More concretely, during the training stage, we

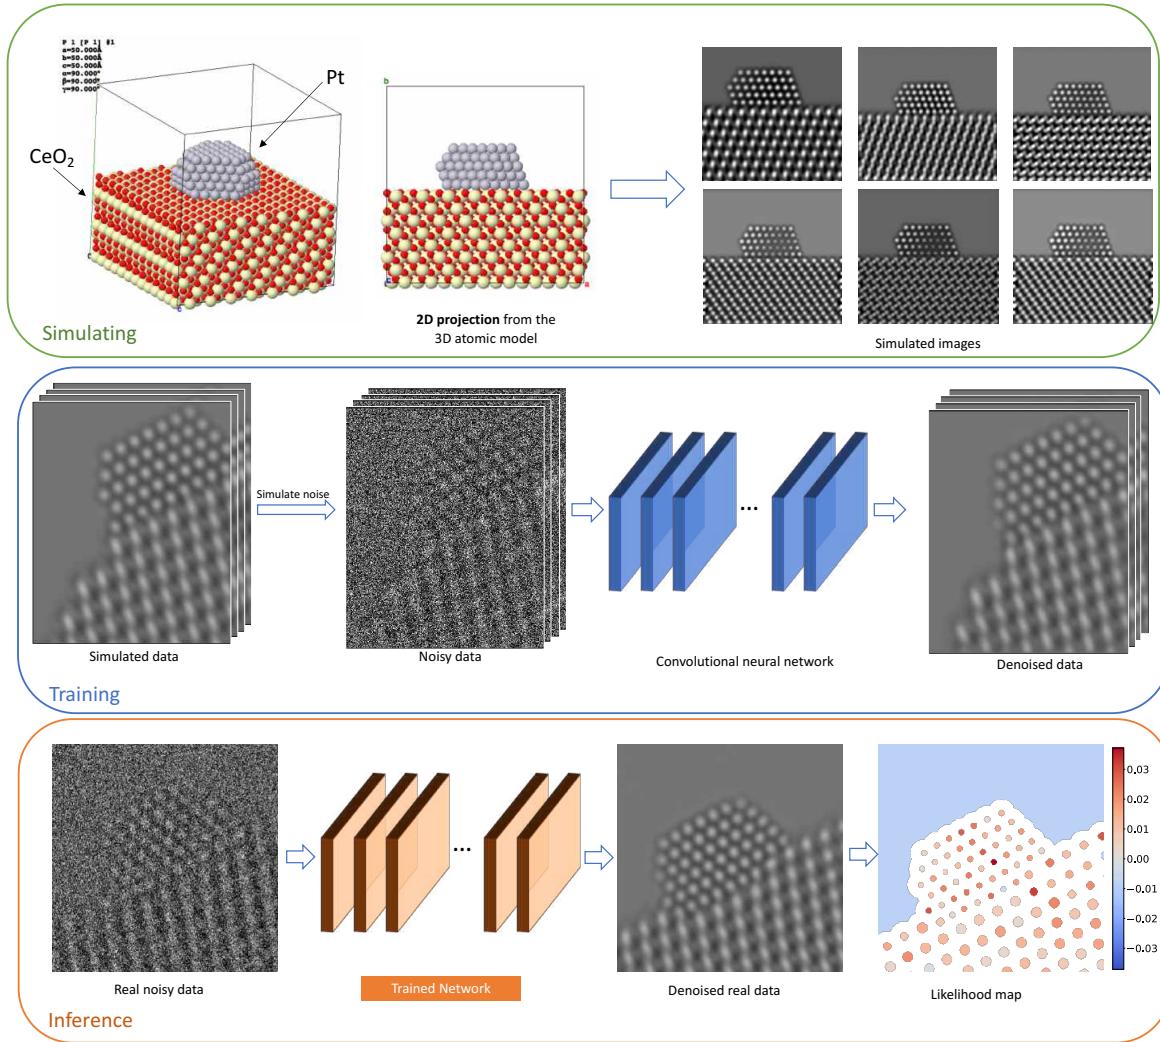


Figure 2.1: **Simulation-based denoising framework.** (Top) A training dataset is generated by simulating TEM images of different structures at varying imaging conditions. Here we focus on structures of Pt nanoparticles supported on CeO_2 . (Middle) A CNN is trained using the simulated images, paired with noisy counterparts obtained by simulating the relevant noise process. (Bottom) The trained CNN is applied to real data to yield a denoised image. After analyzing the image to extract structures of interest, a likelihood map is generated to quantify the agreement between this structure and the noisy data.

compute the parameters by solving

$$(3.1) \quad \hat{\theta} = \arg \min_{\theta} \mathbb{E} \left[\sum_{i=1}^N \mathcal{L}(f_{\theta}(Y(x_i)), x_i) \right] = \arg \min_{\theta} \mathbb{E} \left[\sum_{i=1}^N \|f_{\theta}(Y(x_i)) - x_i\|_2^2 \right]$$

We perform minimization iteratively using a variant of stochastic gradient descent. We approximate the expectation in Equation 3.1 by sampling new realizations of the noisy image $Y(x_i)$ every time we

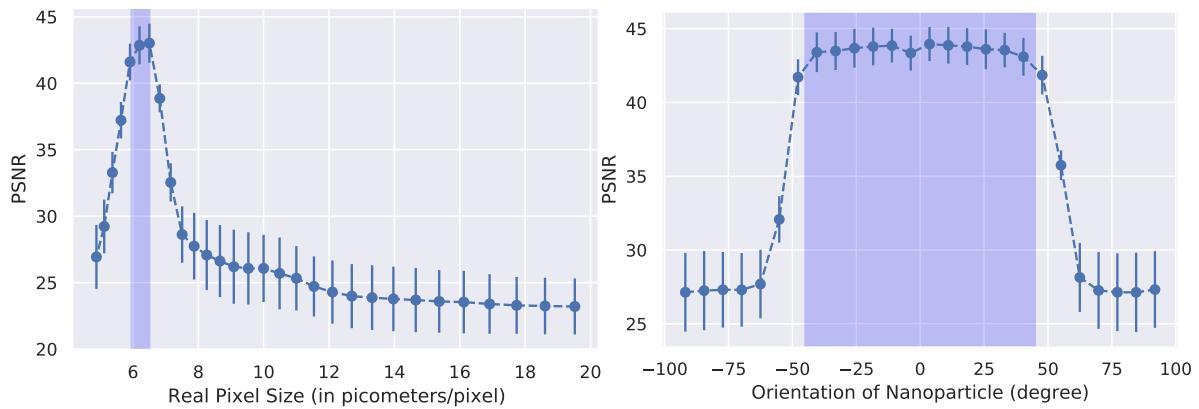


Figure 3.1: **Overfitting scaling and orientation.** The plots show the performance of the proposed network in PSNR for simulated images that have different scalings (on the left, measured in terms of the corresponding pixel size in picometers) and orientations (on the right). The CNN was trained on data augmented with rescaled and rotated images corresponding to the regions shaded in purple. When tested out of those regions, the denoising performance deteriorates significantly. This shows that careful data augmentation is required to ensure invariance to scaling and orientation.

compute the gradient. Once the network is trained, it can be directly applied to new noisy images to perform denoising.

A crucial difference between SBD and previous methodology for deep denoising is that the training set needs to be explicitly *designed*. In order to ensure effective generalization to real data, we must include sufficient variation of imaging parameters and image structure in the training dataset. In addition, particular care is needed to enforce invariance to small changes in the geometry of the image. Figure 3.1 shows that a denoising CNN can easily overfit the specific alignment and scale of the training data. This issue can be addressed by augmenting the training set with rotated and scaled versions of the simulated images. Determining how to optimally sample the space of possible simulation parameters when generating data to train CNNs for denoising is an important methodological question for future research.

3.2. Exploiting non-local signal structure. Images in scientific applications often have pixel-intensity distributions that differ significantly from those of natural images. Our case study shows that it is crucial to take this into account in order to achieve successful denoising. Current state-of-the-art networks for denoising photographic images have very small fields of view. For example, the field of view of DnCNN [75] and DURR [76] are 41×41 pixels, and 45×45 pixels respectively. Unlike photographic images, the TEM images in our case study exhibit very prominent global regularities, due to periodicity in the atomic structure of the imaged materials. In addition, electron-microscopy images are often measured at very low SNRs (in our case, the SNR for the real TEM data is about 3 dB, but most works for photographic images focus on an SNR above 22 dB, see e.g. [75]). As the SNR decreases, denoising CNNs tend to average over larger regions of the surrounding pixels, as demonstrated in Ref. [49] (qualitatively, this is the same behavior observed in a classical linear Wiener filter [73]). These considerations motivate using networks with large field of view to denoise TEM data.

(a) TEM Images					
MODEL	Parameters	FoV	PSNR	SSIM	
SBD + DnCNN [75]	668K	41 × 41	30.47 ± 0.64	0.93 ± 0.01	
SBD + Small UNet [76]	233K	45 × 45	30.87 ± 0.56	0.93 ± 0.01	
SBD + UNet (32 base channels)	352K	221 × 221	36.39 ± 0.77	0.98 ± 0.01	
SBD + UNet (64 base channels)	1.41M	221 × 221	37.24 ± 0.76	0.99 ± 0.01	
SBD + UNet (128 base channels)	5.61M	221 × 221	38.05 ± 0.81	0.99 ± 0.01	
SBD + UNet (128 base channels)	70.15M	893 × 893	42.87 ± 1.45	0.99 ± 0.01	

(b) Photographic Images						
MODEL	Parameters	FoV	PSNR		SSIM	
			$\sigma = 30$	$\sigma = 70$	$\sigma = 30$	$\sigma = 70$
UNet	102K	49 × 49	29.67 ± 2.84	26.16 ± 2.79	0.83 ± 0.06	0.70 ± 0.09
UNet	352K	221 × 221	29.65 ± 2.76	26.08 ± 2.68	0.83 ± 0.05	0.70 ± 0.08
UNet	4.4M	893 × 893	29.54 ± 2.82	26.07 ± 2.80	0.83 ± 0.06	0.70 ± 0.09

Table 3.1: **Field of view of CNN architectures and performance.** Mean PSNR and SSIM (\pm standard deviation) of different CNN architectures on the (a) held-out simulated test set of TEM data described in Section 5.2 and (b) validation set of the DIV2K photographic image dataset [1]. For TEM images, increasing the field of view (FoV) of the UNet from 45×45 pixels to 221×221 produces a dramatic increase of around 6 dB in PSNR, even if the number of parameters remains similar. Increasing the number of parameters while keeping the field of view constant produces a modest gain in performance. In contrast, changing the field of view has almost no effect on denoising photographic images. The networks used for photographic images were trained on 512×512 patches extracted from training images of DIV2K [1] corrupted with additive Gaussian noise with standard deviation $\sigma \in [0, 100]$.

Here we propose to denoise TEM data using deep networks with very large fields of view: 221×221 pixels and 893×893 pixels, a 25-fold and 400-fold increase with respect to generic denoising architectures respectively. In order to obtain a large field of view efficiently (i.e. without dramatically increasing the number of parameters in the network), we propose using a UNet network architecture [61]. We use 4 downsampling operations to achieve the 221×221 field of view and 6 downsampling operations to achieve the 893×893 field of view (see Section B for a detailed description of the architecture). Table 3.1 compares the influence of the field of view in denoising photographic and TEM images. For photographic images the performance of the network remains almost constant as we increase the field of view. In contrast, for TEM images increasing the field of view produces a dramatic improvement in performance (6 dB and 10 dB, when the field of view is 221×221 and 893×893 respectively). Increasing the number of parameters, while keeping the field of view constant, has a very modest effect, which suggests that the increase in field of view is the primary reason for the improvement.

In order to gain some insight into the denoising mechanisms learned by our models, we apply the gradient-based analysis proposed by Ref. [49]. We visualize the linear term in the first-order Taylor decomposition of the denoising map with respect to its input for specific pixels. In more detail, we

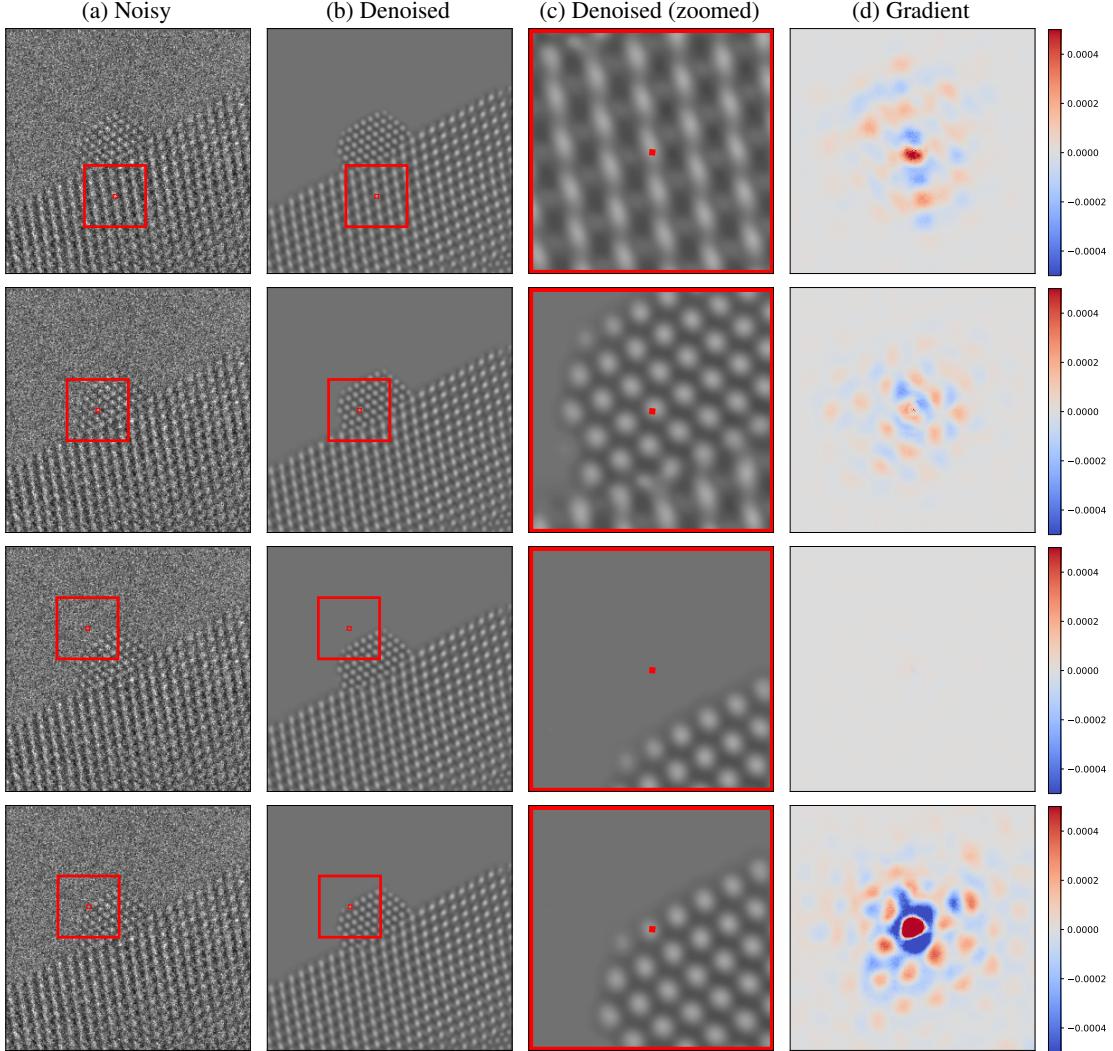


Figure 3.2: Gradient analysis of the learned denoising function on real data. (a) To compute the red pixel in the denoised image (b), the proposed CNN can use noisy pixels in a neighbourhood of up-to 893×893 , of which a 300×300 area (red box) (c) is highlighted. The gradient of the denoised pixel with respect to its input indicates what regions in the noisy image have a greater influence on the estimate (according to a first-order Taylor approximation to the denoising map). The gradient (d) weights nearby pixels more heavily, but also has significant magnitude at pixels located on different atoms. This suggests that the CNN combines local and non-local information to estimate the pixel. The colorbar is shared across all the images.

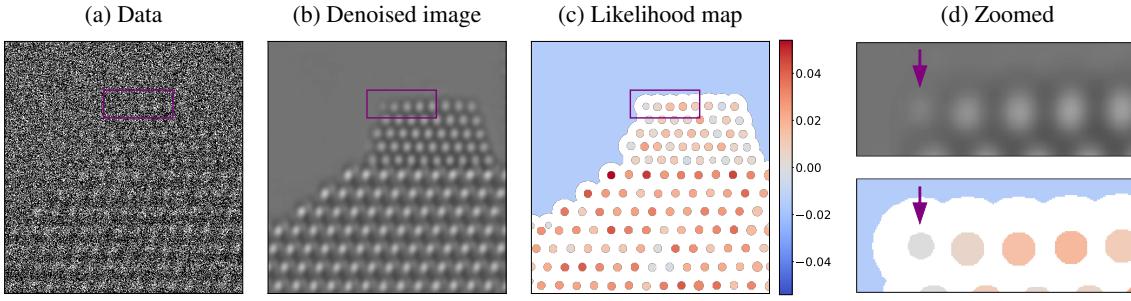


Figure 3.3: **Likelihood map.** When the simulated noisy image in (a) is denoised using the proposed framework (b), a spurious atom appears at the left edge of the nanoparticle (see zoomed image (d)). The value of the likelihood map (c) at that location is very low, indicating that the presence of an atom is less consistent with the observed data than its absence.

compute the gradient of a pixel in the denoised image $f(y)_i$ with respect to the input noisy image y . This vector (or image) $\nabla_y(f(y)_i)$, makes it possible to visualize the influence of different regions of the noisy image on the denoised pixel $f(y)_i$. This approach is similar to visualization methods proposed in the context of image classification (e.g. [63, 50]). Our analysis reveals that the network learns to simultaneously exploit local structure as well as non-local periodicities in the data (see Figure 3.2). This demonstrates the remarkable flexibility of data-driven denoising based on deep learning.

3.3. Likelihood maps. In most applied domains, the goal of denoising is to uncover image structure of scientific interest. In our case study, this corresponds to the location and intensity of projected columns of atoms in a catalytic nanoparticle that is surrounded by a vacuum. Quantifying to what extent such structure is consistent with the observed measurements is therefore of great interest. We propose to achieve this by computing the likelihood of the data *with respect to meaningful features identified in the denoised image*. The general procedure, and its implementation in the case of our case study, are as follows:

1. Identify a region of interest \mathcal{R} . In our case study, this would correspond to an atomic column, located for example via blob detection [39], or to the vacuum.
2. Fit a low-dimensional model to the denoised image within the region of interest. The low-dimensional model provides an estimate of the image value x_i at each pixel location $i \in \mathcal{R}$. In our case, we assume that the intensity of each atomic column and the vacuum are constant, so the estimate is obtained by averaging over all denoised pixels in \mathcal{R} .
3. Compute the likelihood of the noisy data in \mathcal{R} with respect to the estimated pixel values. In our case, the noise is approximately independent and individually distributed (iid) Poisson (see Section 4.3), so the likelihood is given by

$$(3.2) \quad \mathcal{L}(\mathcal{R}) := \prod_{i \in \mathcal{R}} p_{x_i}(y_i),$$

where y_i denotes the noisy value in the i th pixel, and p_{x_i} is a Poisson probability mass function (pmf) with rate parameter x_i . Note that in the low-dimensional model, which assumes constant intensities, x_i is constant for all pixels in \mathcal{R} .

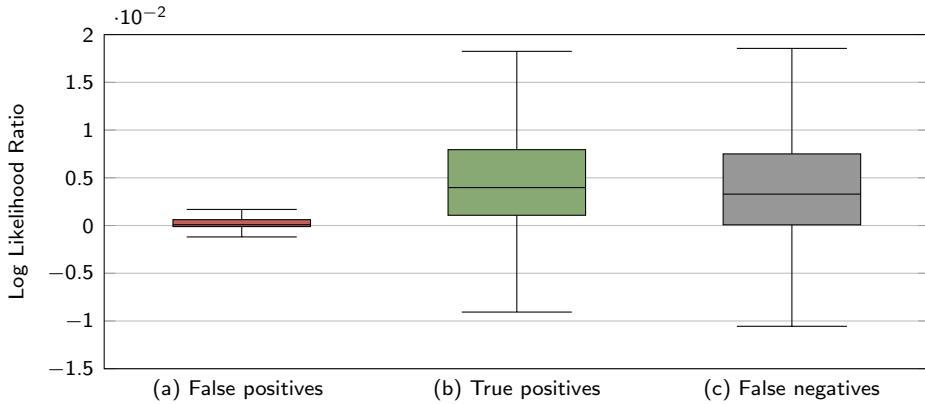


Figure 3.4: **Distribution of likelihood ratio.** The figure shows the distribution of log-likelihood ratio of over 25,000 regions of interest computed from the surface of 1550 denoised images using the dataset described in Section 5.3.2. The empirical distribution is visualized as a box plot indicating the median, 25th quartile, 75th quartile, minimum and maximum value of the distribution. The regions containing spurious atoms (false positives, (a)) have a much lower log-likelihood ratio than the regions containing accurately recovered atoms (true positives, (b)). Regions where existing atoms were not detected (false negatives, (c)) have a higher log-likelihood ratio, comparable to that of the regions with accurately recovered atoms. The occurrence of missing and spurious atoms in denoised images is quite rare: out of the 25,732 regions of interest, only 2,457 and 2,368 were false positives and false negatives respectively.

This technique makes it possible to consider different hypotheses about the underlying image structure and compare their agreement with the observed data. In our case study, we evaluate the hypotheses that a detected atomic column is (1) truly there, or (2) an artefact introduced by the denoising procedure. The likelihood under hypothesis (1) is computed as above. The likelihood under hypothesis (2) is computed by setting the estimate x_i equal to the average intensity of the noisy pixels identified as belonging to the vacuum region. To visualize the consistency of the two hypotheses with the measured data, we plot the difference in their log likelihood for each region of interest. This is equivalent to performing a log-likelihood ratio test. We call this visualization a *likelihood map*.

Figures 1.1 and 3.3 show likelihood maps for the real data and for a simulated example. In the simulated example (Figure 3.3), a spurious atom is detected at the left end of the zoomed region. However, the likelihood map at that location is very low, which indicates that the presence of an atom is not consistent with the observed data at that location. Figure 3.4 shows the distribution of log-likelihood ratio of over 25,000 regions of interest extracted from the surface of over 1,550 denoised images obtained from the dataset described in Section 5.3.2. As shown in Figure 3.4, the log-likelihood ratio values of spurious atoms (false positives, (a)) are much lower than those of correctly-identified atoms (true positives, (c)). When the network fails to detect atoms (false negatives, (b)), we observe that the log-likelihood ratio in such regions tends to be high. It is worth noting that the occurrence of spurious and missing atoms in the denoised images is quite rare: out of the 25,732 regions identified, only 2,457 and 2,368 regions correspond to spurious and missing atoms respectively.

Visualizing the likelihood is useful to quantify the agreement between the output of deep-learning models and the observed data, but it is important to note that the approach suffers from sampling bias. We focus on regions of the input that have been selected because they resemble structures of scientific interest. The data in those regions are therefore more likely to be in agreement with the presence of such structures, just by the sheer fact that they have been selected. This is a manifestation of the notorious multiple-comparisons problem [8, 7]. Overcoming this issue is an important challenge for future research.

4. Dataset. The TEM image data used in this work correspond to images from a widely utilized catalytic system, which consist of platinum (Pt) nanoparticles supported on a larger cerium (IV) oxide (CeO_2) nanoparticle. This bi-functional catalytic system is ubiquitously used in clean energy conversion and environmental remediation applications, in addition to a broad range of other chemical reactions [51, 74, 53]. From a general point of view, this system can be considered as a model for supported nanoparticle catalysts, since a large number of heterogeneous catalysts are based on metallic nanoparticles supported over different oxides. Thus, results and conclusions extracted from the current work are relevant to a great number of similar samples in the field of catalysis (e.g., oxide crystals supporting metal nanoparticles).

4.1. Real Data. The real data used to test the proposed SBD framework consist of a series of images of the Pt/ CeO_2 catalyst. The images were acquired in a N_2 gas atmosphere using an aberration-corrected FEI Titan transmission electron microscope (TEM), operated at 300 kV and coupled with a Gatan K2 IS direct electron detector. The detector was operated in electron counting mode with a time resolution of 0.025 sec/frame and an incident electron dose rate of 5,000 $\text{e}^-/\text{\AA}^2/\text{s}$. The electromagnetic lens system of the microscope was tuned to achieve a highly coherent parallel beam configuration with minimal low-order aberrations (e.g., astigmatism, coma), and a third-order spherical aberration coefficient of approximately $-13\text{ }\mu\text{m}$.

4.2. Simulation Dataset. The simulated TEM image dataset was generated using the multi-slice TEM image simulation method, as implemented in the Dr. Probe software package [4] (see Section A.1 for more details on the simulation process). Images were simulated with 1024 x 1024 pixels and then binned to match the approximate pixel size of the experimentally acquired image series. To equate the intensity range of the simulated images with those acquired experimentally, the intensities of the simulated images were scaled by a factor which equalized the vacuum intensity in a single simulation to the average intensity measured over a large area of the vacuum in a single 0.025 second experimental frame (i.e., 0.45 counts per pixel in the vacuum region).

In the type of phase-contrast TEM imaging performed in this work, multiple electron-optical and specimen parameters can give rise to complex, non-linear modulations of the image contrast. These parameters include the objective lens defocus, the specimen thickness, the orientation of the specimen, and its crystallographic shape/structure. Various combinations of these parameters may cause the contrast of atomic columns in the image to appear as black, white, or an intermediate mixture of the two (see, e.g., Figure A.1). When designing the simulated dataset for the SBD framework, it is necessary to include images simulated under widely varied conditions, in order to cover the breadth of possibilities which may arise during a typical experiment. A skilled microscopist attempts to acquire images under conditions in which the image contrast can be interpreted, which limits the overall size of the parameter space under consideration. However, various instances of defocus, tilt, thickness, and shape/structure

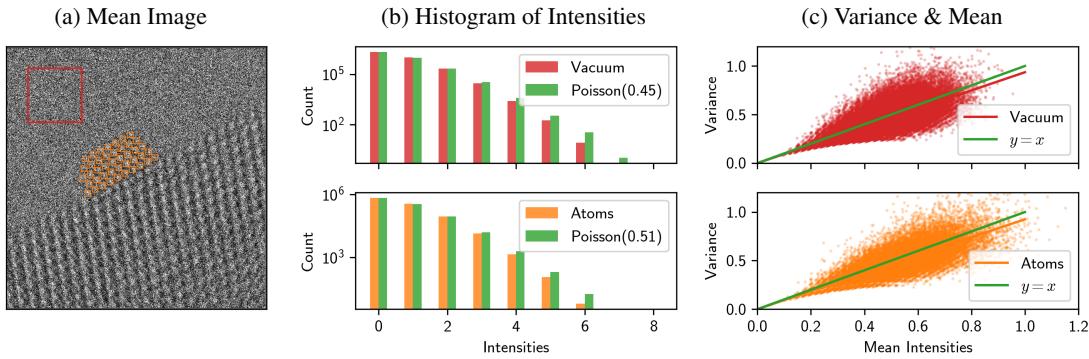


Figure 4.1: **Analysis of the noise in the real data.** The analysis shows that the noise is approximately iid Poisson. (a) Pixel-wise mean over 40 frames of the real data described in Section 4. (b) Histogram of noisy pixel intensities from highlighted regions in the image compared to a simulated Poisson distribution. (c) The plot of empirical mean and standard deviation of pixels approximately follows a line with unit slope, as expected from iid Poisson samples (the spread is due to averaging over only 40 frames).

inevitably arise. To generate our dataset we systematically varied these parameters to produce a large number of potential combinations (approximately 18,000), as described in Sections A.2 and A.3.

4.3. Noise model. The TEM images were acquired on a direct electron detector operating in electron counting mode. In such conditions, the electron dose rate per pixel is sufficiently low enough that individual electron arrivals can be detected and registered. It is well known that the statistical fluctuations of such counting processes for discrete events are governed by shot noise, which can be modeled with a Poisson distribution [3]. We expect that other sources of noise, including fixed pattern noise, dark noise, and thermal noise are minimal after applying a gain correction and a dark reference to the raw image, and by cooling the detector to -20 °C, respectively.

We empirically verified that the noise follows Poisson statistics through the analysis shown in Figure 4.1. Our real dataset, described in Section 4.1 consists of 40 noisy frames acquired sequentially across time in 0.025 sec intervals. Figure 4.1(a) shows the mean image over all 40 frames. The region of the image containing no material (red box) corresponds to the vacuum, where the electron beam intensity is uniformly illuminating the detector. Fluctuations of the intensity in this region therefore purely arise from the noise. We validate that the histogram of these pixels aggregated over the identified spatial region closely follows a Poisson distribution (Figure 4.1(b)). Pixels aggregated over spatial domains corresponding to the Pt atomic columns (orange boxes) show similar behavior. Further, if the distribution is indeed Poisson, the mean and variance of the noisy pixels should be approximately the same. The empirical mean and variance, computed by averaging over the 40 time frames at every spatial location, follows a linear trend, thus further confirming that the noise distribution has Poisson properties (Figure 4.1(c)). The spread in the scatter plot is due to the limited number of time frames over which we average.

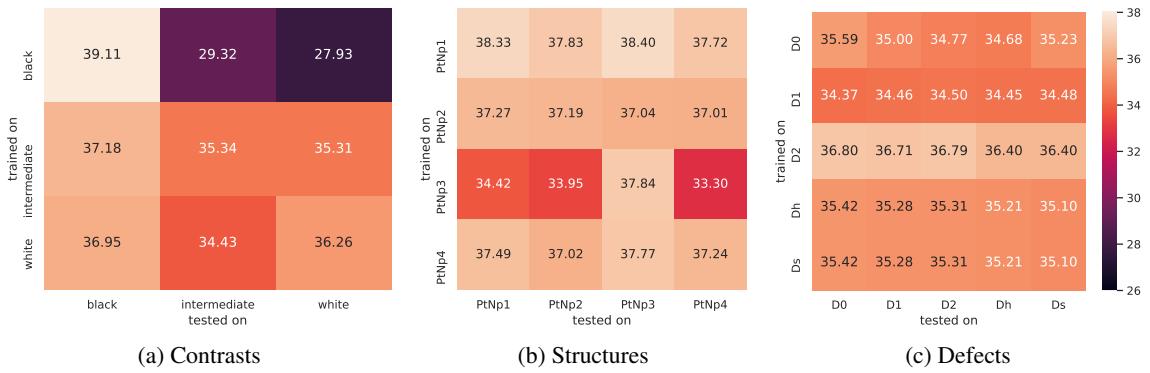


Figure 5.1: **Generalization across different imaging parameters and signal structures.** In order to study the generalization ability of the proposed method we divided the simulated dataset described in Section 4 into subsets, based on the atomic column contrast, the structure/size of the supported Pt nanoparticle, and the defects of the Pt surface structure. The tables show the test PSNR for networks trained and tested on different combinations of the subsets. (a) Networks trained on white and intermediate contrast generalize well to all other contrasts. The network trained on black contrast does not generalize as well. (b) Networks trained on a type of nanoparticle structure generalize well to all other types. (c) Networks trained on one type of defect or no defects generalize well to different types. Figures A.1, A.5, A.4 and A.2 show the effect of these parameters on the images.

5. Experiments and Results. In this section, we evaluate the performance of our proposed methodology and show that we outperform other methods by a large margin (more than 12 dB in PSNR on held-out simulated data). We also perform a thorough analysis of the generalization capability of our models, demonstrating that the CNNs are robust to variations in imaging parameters and in underlying signal structure. Furthermore, we demonstrate that standard performance metrics for photographs, such as peak signal-to-noise ratio (PSNR) and SSIM [71], may fail to produce a scientifically-meaningful evaluation of the denoising results, and we propose a few alternative metrics to remedy this. Finally, we show that our approach achieves effective denoising of real experimental data.

We use CNNs with the proposed UNet architecture with 128 base channels and 6 scales in all of our experiments (see Sections 3.2 and B for more details). The networks were trained on 400×400 patches extracted from the training images and augmented with horizontal flipping, vertical flipping, random rotations between -45° and $+45^\circ$, and random resizing by a factor of 0.75-0.82. The models were trained using the Adam optimizer [28], with a default starting learning rate of 10^{-3} , which was reduced by a factor of 2 every time the validation PSNR plateaued. Training was terminated via early stopping based on validation PSNR. The details of training, validation and test data for each experiment are provided in the corresponding section. Since the models are trained on 400×400 patches, when applying them to larger images we divide the images into overlapping 400×400 patches, denoise them, and then combine them via averaging.

5.1. Generalization to unseen structures and acquisition conditions. In order to study the generalization ability of the proposed approach across different imaging parameters and signal structures we divided the simulated dataset described in Section 4 into different subsets. These subsets

METHODS	PSNR	SSIM
Raw	3.56 ± 0.03	0.00 ± 0.00
Low Pass Filter [52]	21.59 ± 0.07	0.44 ± 0.03
Adaptive Wiener Filter [38]	22.42 ± 1.08	0.63 ± 0.02
VST + NLM [10]	26.55 ± 0.16	0.73 ± 0.01
VST + BM3D [43]	25.27 ± 0.15	0.80 ± 0.01
PURE-LET [41]	28.36 ± 0.88	0.93 ± 0.01
SBD + DnCNN [75]	30.47 ± 0.64	0.93 ± 0.01
SBD + Small UNet [76]	30.87 ± 0.56	0.93 ± 0.01
SBD + Proposed Architecture	42.87 ± 1.45	0.99 ± 0.01

Table 5.1: **Results on simulated test data.** Mean PSNR and SSIM (\pm standard deviation) of different denoising methods on the held-out simulated test set described in Section 5.2. SBD approaches achieve the best results. SBD combined with the proposed architecture outperforms all other techniques by about 12 dB. The performance of SBD applied to additional architectures is reported in Table 3.1.

were classified based on (1) the character of the atomic column contrast, (2) the structure/size of the supported Pt nanoparticle, and (3) the defects of the Pt surface structure. The contrast was classified into three divisions, black, intermediate, or white contrast, by a domain expert (see Figure A.1 in the supplementary material). The nanoparticle structure was classified into four categories, “PtNp1” through “PtNp4”. PtNp1 and PtNp2 correspond to supported Pt nanoparticles of size 2 nm, which differ in the presence or absence of an atomic column located at the interface between the Pt and the CeO_2 support. PtNp3 corresponds to a Pt nanoparticle 1 nm in size. PtNp4 corresponds to a Pt nanoparticle 3 nm in size. Finally, the defects were divided into five categories: “D0”, “D1”, “D2”, “Dh”, and “Ds” in accordance with the atomic-scale structural models presented in A.1 and in particular in Figure A.5. D0 is the initial structure, D1/D2 a structure in which 1/2 atomic columns have been removed respectively, Dh a structure in which a column has been reduced to half its original occupancy, and Ds a structure in which a column has been reduced to a single atom. The generalization ability of the proposed CNN was evaluated by systematically training on each of the subsets and testing on the rest. The number of images in each subset was fixed to be equal in order to ensure a fair comparison.

The performance of SBD is robust to variations in imaging parameters and in the underlying signal structure, as shown in Figure 5.1. We only observe a significant decrease in performance when the network is trained on black-contrast images and tested on other contrasts (interestingly the network generalizes well from white and intermediate contrasts to black contrasts).

5.2. Comparison of SBD with other methods. The imaging parameters of the real data, described in Section 4, are well described by the white contrast category defined in Section A.2. We therefore used the subset of simulated dataset corresponding to this contrast (5583 images) to compare our proposed methodology to other models. 90% of the data were used for training. The remaining 559 images were evenly split into validation and test sets. We compare our proposed UNet architecture (see Section B) with two state-of-the-art architectures for photographic-image denoising [75, 76] (see Sections B.2 and B.3), and with several classical denoising methods:

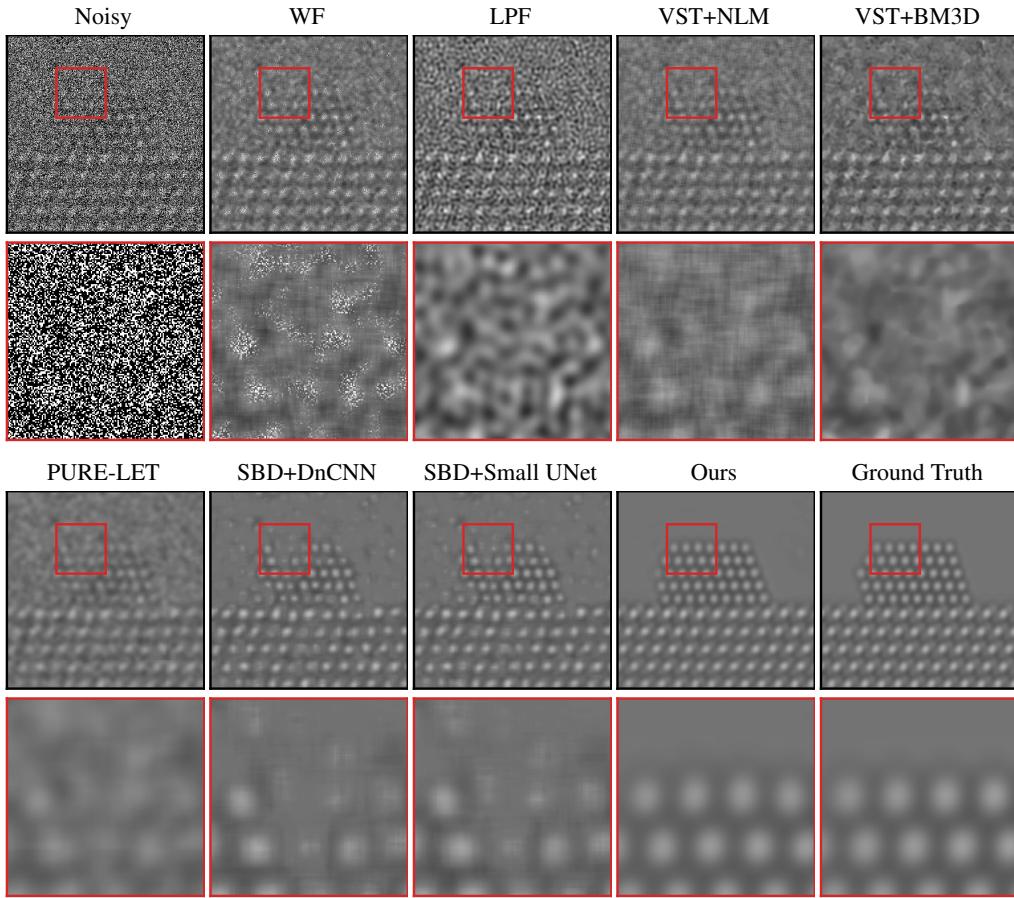


Figure 5.2: **Denoising results for simulated data.** Comparison of SBD and the baseline methods described in Sections 5.2 and B. The second row zooms in on the region in red box. Our proposed approach produces images of much higher quality than the other approaches, and is able to accurately recover the atomic structure of the nanoparticle. For example, the vacuum region in images denoised by several of the baselines contain visible artefacts, including missing atoms. See Figure C.1 in the supplementary material for an additional example.

- **Low-pass Filter (LPF):** We apply a standard low-pass filter [52] designed by a domain expert. The cut-off frequency was determined based on inspection of the data in the Fourier domain.
- **Adaptive Wiener Filter:** We apply an adaptive low-pass Wiener filter [38] to account for variations in local image statistics. The mean and variance around each pixel are estimated from a local neighborhood. We selected a neighborhood with radius 13 pixels based on expert evaluation of the denoised images in the validation set.
- **VST + BM3D and VST + Non-Local Means (NLM):** BM3D [43] and NLM [10] are popular denoising techniques for natural images with additive Gaussian noise, which can be adapted to Poisson noise by applying a nonlinear variance-stabilizing transformation (VST) [77]. More

specifically, we use the Anscombe transformation proposed in Ref. [43].

- **PURE-LET**: PURE-LET [41] is a transform-domain thresholding algorithm adapted to mixed Poisson–Gaussian noise. The method requires the input image to have dimensions of the form $(2^k, 2^k)$. To apply this method on our TEM images, we extracted 128×128 overlapping patches, denoised them and combined them via averaging.

For all methods, hyperparameters were chosen based on the validation data. Performance was measured in terms of SSIM [71] and peak signal-to-noise ratio (PSNR).

The results demonstrate that SBD is an effective denoising methodology for TEM data. Our proposed CNN outperforms all other methods by a margin of 12 dB in PSNR on the simulated test data, as shown in Table 5.1, and Figures 5.2 and C.1. SBD recovers the overall shape of the nanoparticle, the interface between the nanoparticle and the support, and the different periodic patterns of the CeO_2 support and Pt nanoparticle. Contrast features, such as subtle patterns of bright, intermediate and dark features associated with the atomic structure of the CeO_2 crystal, are well reproduced in the images denoised via SBD, but are mostly absent from the results of the baseline approaches.

5.3. Beyond PSNR: Towards scientifically-meaningful evaluation metrics. Domain scientists denoise images in order to extract scientifically relevant information. In our case, the atoms on the surface of nanoparticles are of particular interest, because the atomic configuration at the surface regulates the nanoparticle’s ability to catalyze chemical reactions. It is therefore of critical importance to understand how different denoising methods recover these atoms. We can verify visually that SBD achieves a largely successful recovery in held-out simulated data, whereas the baseline methods described in Section 5.2 do not (see Figure 5.2 for example). However, visual inspection is a limited and non-quantitative evaluation tool. Unfortunately, standard metrics like PSNR and SSIM are insensitive to changes in the atomic structure of the nanoparticle surface, because these changes have a small effect on the overall intensity of the images. We demonstrate the lack of sensitivity through a synthetic example in Figure 5.3: when we add or remove an atom in the surface the PSNR and SSIM remain roughly constant. Motivated by the need for scientifically-relevant performance evaluation, we propose several metrics explicitly designed to account for changes in surface atomic configuration in Section 5.3.1. We report an evaluation of SBD using these metrics on a challenging test dataset in Section 5.3.2.

5.3.1. Evaluation metrics. To define metrics that evaluate detection of surface atoms, we assume that there is a predefined approach to perform detection based on the denoised images. In our case of interest, we apply a blob detection algorithm (Laplacian of Gaussian [39]) to locate the centers, and compute the α -shape of all the atom centers using Delaunay triangulation [57]. Let A and B be the set of surface atoms of interest in the denoised image and the ground-truth clean image respectively. We propose the following four metrics to measure the fidelity of the recovered structure:

- **Precision** is the fraction of atoms in the denoised image that are also present in the clean image.

$$(5.1) \quad P(A, B) = \frac{|A \cap B|}{|B|}.$$

- **Recall** is the fraction of atoms in the clean image that are correctly recovered in the denoised image.

$$(5.2) \quad R(A, B) = \frac{|A \cap B|}{|A|}.$$

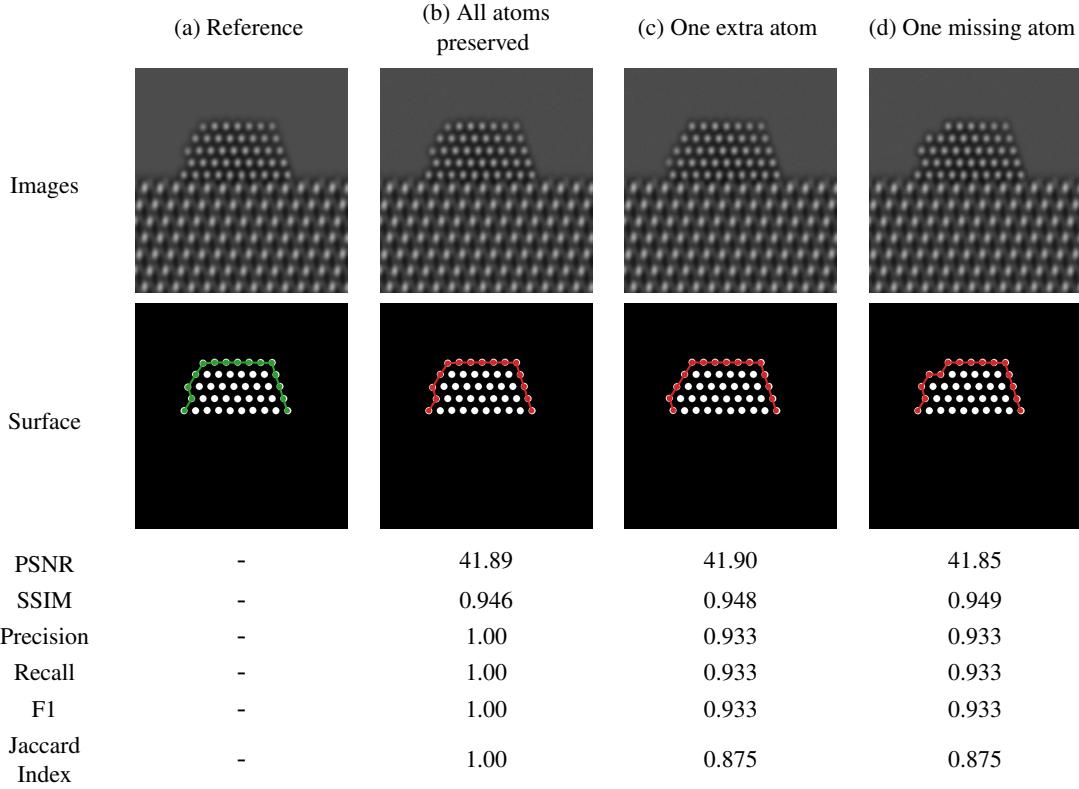


Figure 5.3: **Scientifically-meaningful metrics for atom detection.** In order to compare metrics in terms of their sensitivity to changes in atomic structure we perturb and compare the reference image in (a) to several modified images. In (b) the atomic structure is the same. In (c) a spurious atom is added to the image in the top left. In (d) an atom is removed from the top left. The images in (b), (c), and (d) are further corrupted by adding iid Gaussian noise with a small deviation of around 2/255. The PSNR and SSIM [72] of (b), (c) and (d) with respect to (a) are essentially constant, indicating that these metrics are not sensitive to changes in atomic configuration. In contrast, our proposed metrics reflect these changes more accurately: (b) is assigned a score of 1 in all metrics, whereas (c) and (d) are consistently assigned lower values.

- **F1 score** combines precision and recall by giving them equal importance.

$$(5.3) \quad F(A, B) = 2 \frac{P(A, B)R(A, B)}{P(A, B) + R(A, B)}.$$

- **Jaccard index** is an alternative measure consisting of the ratio between the size of the intersection between the recovered atoms and the ground truth divided by the size of their union.

$$(5.4) \quad J(A, B) = \frac{|A \cap B|}{|A \cup B|}.$$

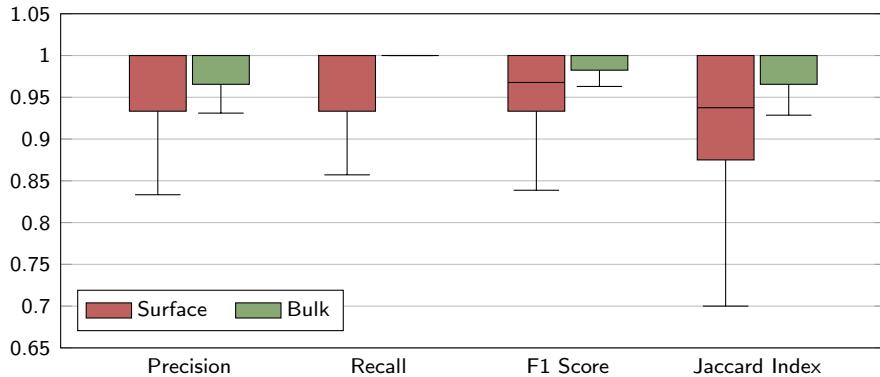


Figure 5.4: Performance of SBD in terms of our proposed metrics. We compute all our proposed metrics (see Section 5.3.1) on over 7,000 denoised images corresponding to 25 unique noisy images sampled from the 308 clean images described in Section 5.3.2. The empirical distribution on the surface (red) and bulk (green) is visualized as box plots indicating the median, 25th quartile, 75th quartile, minimum and maximum value of the distribution. SBD has a near perfect performance in the bulk with all metric values hovering around 1. On the surface, SBD achieves a median score of 1 for precision and recall, and about 0.95 for F1 score and Jaccard index.

When performing intersection and union operations, we consider two atoms to be the same if the distance between their centers is less than a threshold of 100 pixels. All our metrics take values between 0 and 1 (1 is best). Figure 5.3 shows a synthetic example comparing three images with different atomic configurations: all the images have similar PSNR and SSIM values, but the precision, recall, F1 and Jaccard index show substantial differences.

5.3.2. Evaluating atom detection accuracy. To evaluate the performance of the proposed approach to recover atoms at the surface, we designed a new dataset with 308 images, where the imaging parameters are set based on the real dataset described in Section 4.1. This new dataset is similar to the one used for the generalization experiments in Section 5.1, but here we add more diverse surface defects. We created a series of 44 Pt/CeO₂ structural models with atomic-level surface defects such as the removal of an atom from a column, removal of two atoms, removal of all but one atom and addition of an atom at a new site (see Figure C.2 for a visual overview). We hypothesize that these defects emulate dynamic atomic-level reconfigurations that could potentially be observed in real experiments. To match the image contrast of our real data, we simulated images under defocus values ranging from 6 nm to 10 nm, all with a tilt of 3° in x and -1° in y and a support thickness of 40 Å. SBD recovers all the atoms in the bulk almost perfectly, as reflected in the different metrics. On the surface, SBD achieves a median score of 1 for precision and recall, and more than 0.95 on F1 and Jaccard index (see Figure 5.4).

5.4. Performance on real data. In the experiments reported in Sections 5.2 and 5.3 we used a network trained on all simulated images from the white contrast category defined in Section 5.1. However, the real data described in Section 4 more closely corresponds to a subset of white contrast images satisfying the following conditions: structure limited to PtNP2, thickness between 40 Å - 60 Å

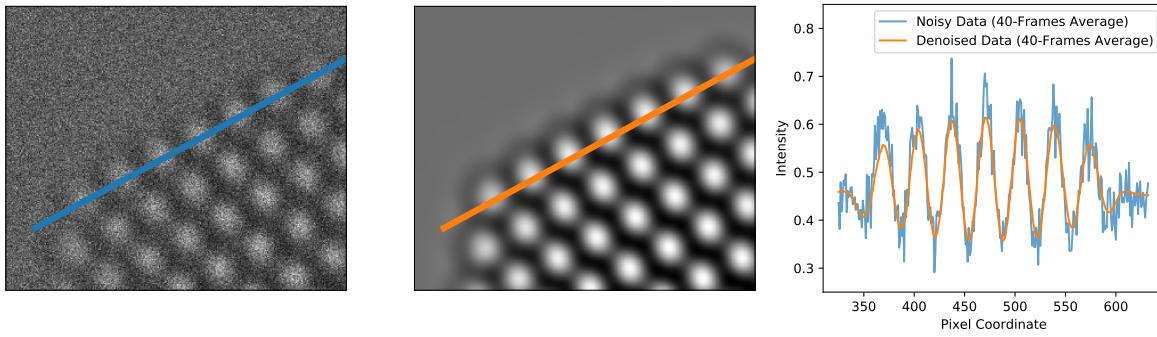


Figure 5.5: **Validation on real data.** The real data consist of 40 frames which are approximately stationary and aligned. Their temporal average (left) therefore provides a reasonable estimate for the true intensity profile. In the image on the right, we compare the average intensity profile on the surface atomic columns of the platinum nanoparticle for the denoised data (middle) and the temporal average (left). The profiles are very similar (except for some spurious fluctuations in the temporal average), which suggests that the proposed approach achieves effective denoising on the real data.

and, defocus between 5 nm and 10 nm. We used 236 images from this subset for training, and another such 15 images for validation. We also trained two state-of-the-art architectures for photographic image denoising - DnCNN [75] and DURR [76] on these data.

Results on real experimental data obtained using SBD trained on this relevant subset of white contrast are shown in Figures 1.1, 5.6, and C.3. SBD produces denoised images that are of much higher quality than those of the baseline methods described in Section 5.2, which contain obvious artefacts. Further, we validate the denoising results of SBD by comparing to an estimated reference image obtained by temporal averaging. Our real dataset consists of 40 frames that are approximately stationary and aligned. Therefore, their temporal average provides a good estimate for the ground-truth images. As shown in Figure 5.5, the denoised intensity values of the atomic column approximately match those of the estimated reference image.

In the rest of this section, we compare the performance of SBD and unsupervised denoising techniques on the real experimental data, and analyze the effect of the design of the training dataset on the denoised output produced by SBD.

5.4.1. Comparison to unsupervised deep denoising methods. Unsupervised denoising techniques can be used to train a denoising CNN using only noisy images (see Section 2 for a discussion on this methodology). We apply the following unsupervised methods to the real data described in Section 4.1:

- **Blind-spot net** [32] is a CNN which is constrained to predict the intensity of a pixel as a function of the noisy pixels in its neighbourhood, without using the pixel itself. This method is competitive with the current supervised state-of-the-art CNN on photographic images. However, when applied to our real dataset it produces denoised images with visible artefacts (see Figure 5.7(c)). A possible explanation is the limited amount of data (40 noisy images) we train on. To validate this hypothesis, we trained a blind-spot net on simulated training sets of different sizes. The performance on held-out data is indeed poor when the training set is

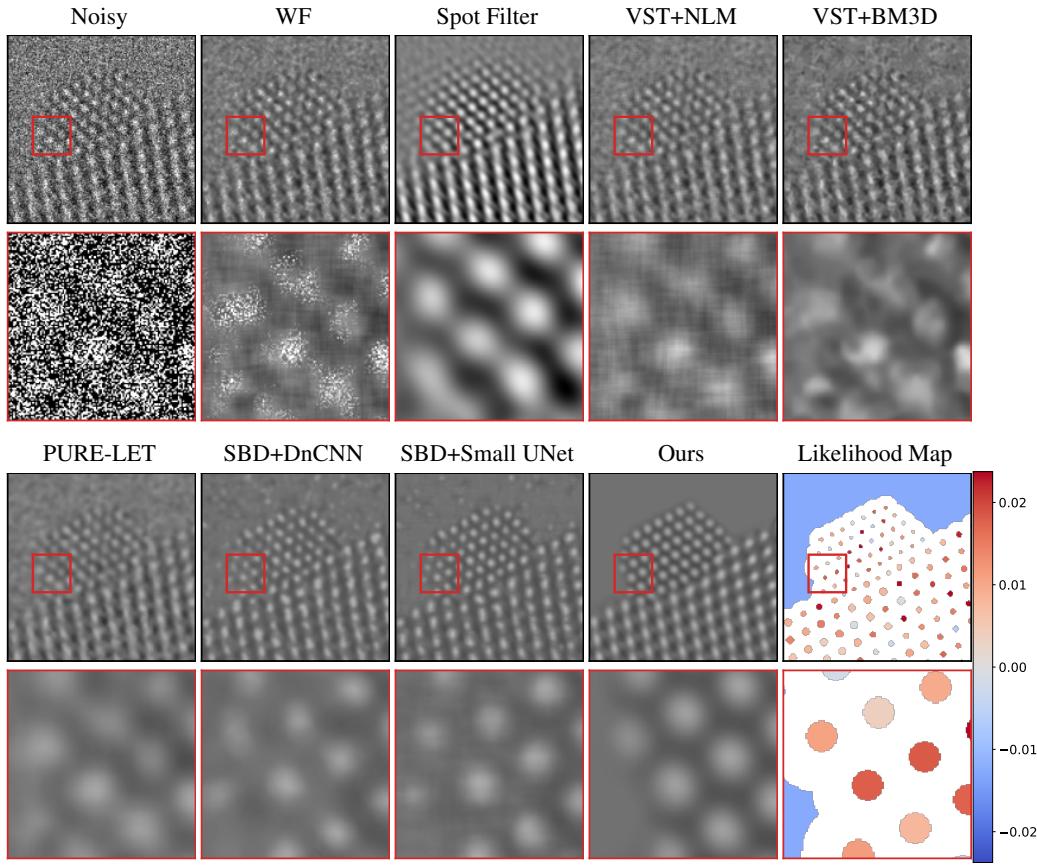


Figure 5.6: Denoising results for real data. Comparison SBD and the baseline methods described in Sections 5.2 and B when applied on the real data described in Section 4.1. The second row zooms in on the region in red box. In contrast to the other methods, SBD combined with the proposed architecture is able to precisely recover the structure of the nanoparticle and has very few artefacts, particularly in the vacuum region. The likelihood map quantifies the agreement between recovered structures in the denoised images, such as atomic columns and the vacuum, and the observed data (see Section 3.3 for more details). See Figure C.3 for an additional example

small, but it improves to the level of supervised approaches as we use more training data (see Figure 5.8).

- **Blind-spot net with early stopping.** In Ref. [62] it is shown that early stopping based on noisy held-out data can boost the performance of blind-spot nets. Here we used 35 images for training the blind-spot net and the remaining 5 images as a held-out validation set. We chose the model parameters that minimized the mean squared error between the noisy validation images and the corresponding denoised estimates. The results (shown in Figure 5.7(d)) are significantly better than those of the standard blind-spot network. However, there are still noticeable artefacts, which include missing atoms.

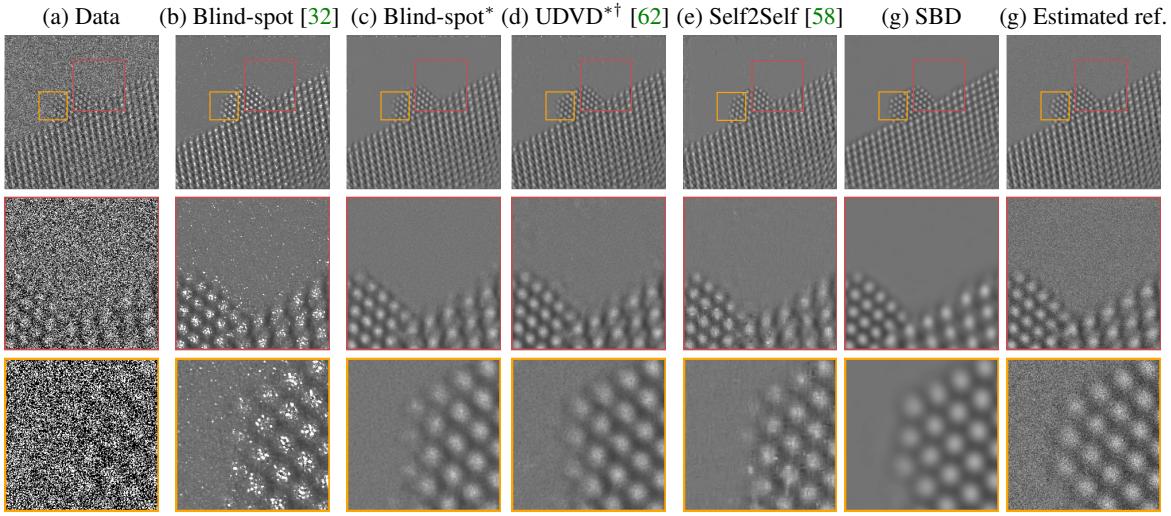


Figure 5.7: **Comparison of unsupervised denoising methods with SBD on real data.** The real data described in Section 4.1 denoised using SBD and unsupervised methods described in Section 5.4.1. The second and third rows zoom in on the region in red and green boxes respectively. Our proposed method denoises the real data more effectively than the unsupervised approaches. SBD is able to precisely recover the structure of the nanoparticle and has very few artefacts (compare visually to the estimated reference image obtained via time averaging; there are three missing atoms for most unsupervised methods in the third row). A * indicates that the method used early stopping and † indicates that the method uses 5 noisy frames as input.

- **Unsupervised Deep Video Denoising (UDVD) [62]** is an unsupervised method for denoising video data based on the blind-spot approach. It estimates a denoised frame using 5 consecutive noisy frames around it. Our real data consists of 40 frames acquired sequentially. UDVD produces better results than blind-spot net, but still contains visible artefacts, including missing atoms (see Figure 5.7(d)). Note that, UDVD uses 5 noisy images as input, and thus has more context to perform denoising than the other methods (including SBD).
- **Self2Self [58]** is an unsupervised method specifically designed for denoising based on a single noisy image. This approach achieves near state-of-the-art performance in noisy photographic images corrupted with moderate amounts of noise [58]. However, when applied to our data, Self2Self produces images with clear artefacts; some of the atoms are missing and the shape of atoms are distorted (see Figure 5.7(e)).

It is important to note that the backbone architectures of all these methods are UNets with large fields of view, like the one used for SBD. In our experiments, we trained the blind-spot nets and UDVD on 600×600 patches extracted from the real data. We used Adam optimizer [28] with a starting learning rate of 1×10^{-4} which was reduced in half for every 2000 epochs. We trained for a total of 5000 epochs. When performing early stopping, we picked the checkpoint with the best mean squared error on the validation set. Following Ref. [58], Self2Self was trained for 150,000 steps with the Adam optimizer and a starting learning rate of 10^{-4} .

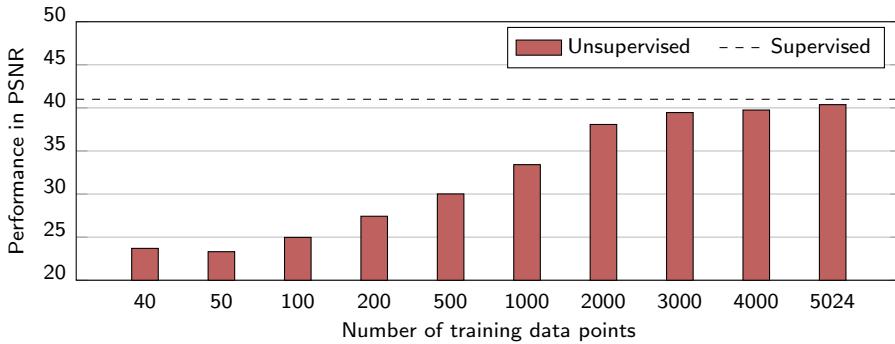


Figure 5.8: **Training set size and unsupervised denoising.** Performance of blind-spot net [32] on a held-out set of images, measured in PSNR, when trained on simulated training sets of different sizes. The held-out set and the training data comprise white contrast images. The performance is poor when the training set is small, but improves to the level of supervised approaches (denoted by dashed line) when trained using more training data. Our real dataset described in Section 4.1 contains only 40 images, so it is likely that the limited amount of data explains the poor performance of the blind-spot net in Figure 5.7.

As shown in Figure 5.7, the unsupervised denoising methods produce higher-quality reconstructions than those of the baseline methods discussed in Section 5.2 (see Figure 5.6). However, they still suffer from visible artefacts, particularly on the surface of the nanoparticle, limiting their practical utility. UDVD is the method that achieves best performance, but it requires multiple noisy frames as input. In contrast, SBD can denoise the image effectively from a single noisy input frame (see Figure 5.7(g)), as long as the simulated training data correspond closely to the real noisy image. Using a single frame is important in some applications, such as our case of interest, where the ultimate goal is to identify dynamic changes in the atomic structure of the nanoparticle.

5.4.2. A word of caution: Effect of training data on SBD. Figures 5.6, C.3 and 5.7 show that SBD achieves impressive results on real data, but it is important to point out that this requires a careful design of the training dataset. Our real data broadly corresponds to images in the white contrast category, defined in Section 5.1. However, when a network trained on white contrast images (Section 5.2) is evaluated on the real data, it produces unnatural streak patterns in the bulk (see third row in Figure 5.9). When visually comparing this to the pattern in the bulk of the reference image computed by time averaging, it is evident that this is an artefact of denoising. This can be remedied by training the network on the more restricted subset of images described in Section 5.4 (see third row in Figure 5.9), whose imaging parameters are more suited to the real acquisition conditions. Since unsupervised denoising methods directly train on the real data, they do not suffer from this problem of mismatch between training and test data. The patterns recovered by unsupervised methods in the bulk are close to the estimated reference image (see Figure 5.9). However, as discussed in Section 5.4.1, they show significant artefacts on the surface of nanoparticle.

6. Discussion and Conclusions. Our case study is a proof of concept that CNNs trained on simulated data can be remarkably effective when applied to real imaging data. It provides several

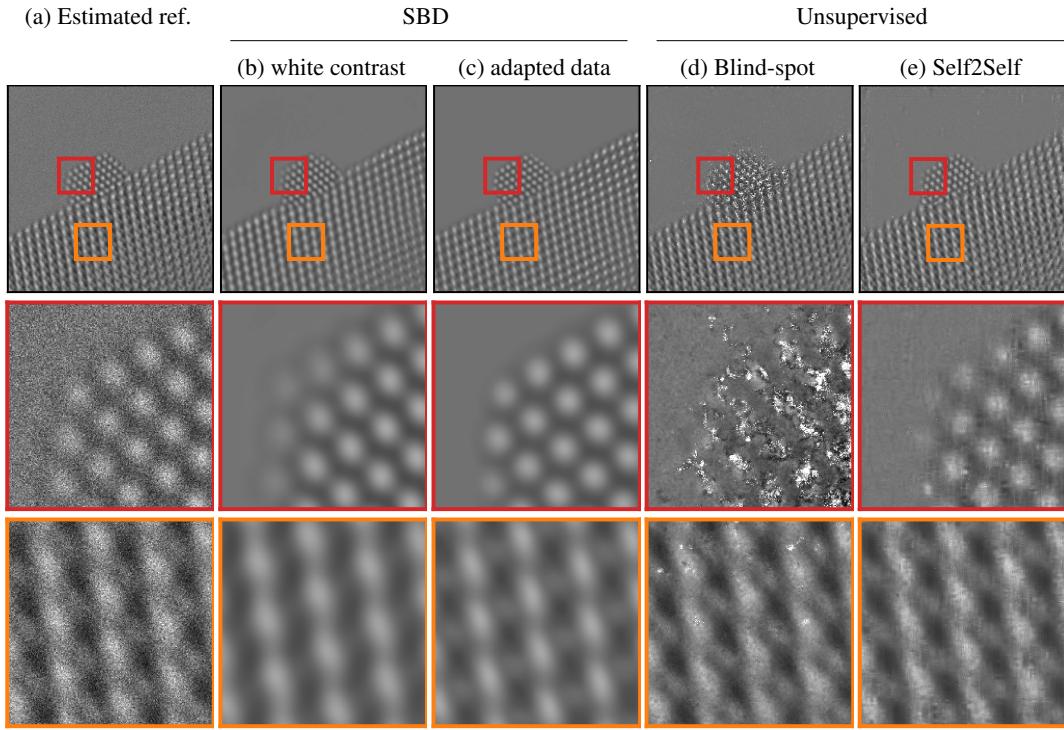


Figure 5.9: Effect of training data on SBD. The real data consists of 40 frames which are approximately stationary and hence their temporal mean (a) can be used as an estimate for the reference image. The second and third row zooms into the red and yellow region respectively. When the network is trained on white contrast (i.e. the training data does not align well with real data), the denoised image of the real data (b) shows an unnatural streak pattern in the bulk (compare row 3 in (b) and (a)). Interestingly, when the training data is more reflective of the real data, the patterns in the bulk are recovered better (row 3 in (c) and (a)). Further, note that unsupervised methods denoise the contrast patterns in the bulk well (row 3 in (d), (e) and (a)) but they suffer from significant artefacts on the nanoparticles (row 2 in (d), (e) and (a)).

insights and suggests future research directions that are relevant, beyond electron microscopy, to other domains where the images of interest can be simulated, such as medical imaging [27, 48], other types of microscopy [19, 69], or astronomy [54]. We show that the design of the training dataset is critical, so an important question is how to design simulated training datasets in a principled systematic way. Answering it will require a deeper understanding of the generalization ability of CNNs with respect to variations in the statistics of the input images. We also demonstrate that architectures tailored to photographic imaging can perform poorly when applied to other data. Designing CNNs for other domains requires an understanding of the image features that are exploited for denoising. Gradient visualization is shown to be useful here, but more advanced visualization techniques are needed. In addition, we demonstrate that standard metrics used to quantify performance in photographs may not be sensitive to scientifically relevant features, and propose several new metrics to address this problem.

Although SBD outperforms other methods by a large margin, some artefacts such as phantom atoms still appear. Our proposed likelihood maps help to flag such events, but may still fail to do so in regions of unusually low SNR. Developing more sophisticated methods for uncertainty quantification is therefore a key research direction. It would also be of great interest to develop unsupervised or self-supervised denoising approaches that are effective with small amounts of data at low SNRs. Finally, to encourage further development of deep-learning methodologies for scientific imaging, we release a denoising benchmark dataset of TEM images, containing 18,000 examples.

Acknowledgements. We gratefully acknowledge financial support from the National Science Foundation (NSF). NSF NRT HDR Award 1922658 partially supported SM who designed, implemented and analyzed the proposed methodology. NSF CBET 1604971 supported JLV and PAC who acquired and processed the experimental data. NSF OAC-1940263 supported RM and PAC who created the training dataset. NSF OAC-1940124 and CCF-1934985 supported BT, who contributed to implement the methodology and conduct the experiments, and DM, who supervised the design of the methodology. NSF OAC-1940097 supported CFG who supervised the design, implementation and analysis of the methodology. All authors participated in writing and reviewing the manuscript. The authors acknowledge ASU Research Computing and NYU HPC for providing high performance computing resources, and the John M. Cowley Center for High Resolution Electron Microscopy at Arizona State University.

REFERENCES

- [1] E. AGUSTSSON AND R. TIMOFTE, *Ntire 2017 challenge on single image super-resolution: Dataset and study*, in The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, July 2017.
- [2] X.-C. BAI, G. McMULLAN, AND S. H. SCHERES, *How cryo-em is revolutionizing structural biology*, Trends in Biochemical Sciences, 40 (2015), pp. 49 – 57, <https://doi.org/https://doi.org/10.1016/j.tibs.2014.10.005>, <http://www.sciencedirect.com/science/article/pii/S096800041400187X>.
- [3] N. C. BARFORD, *Experimental measurements: precision, error and truth*, (1967).
- [4] J. BARTHEL, *Dr. probe: A software for high-resolution stem image simulation*, Ultramicroscopy, 193 (2018), pp. 1–11.
- [5] J. BATSON AND L. ROYER, *Noise2self: Blind denoising by self-supervision*, arXiv preprint arXiv:1901.11365, (2019).
- [6] S. BECKOUCHE, J.-L. STARCK, AND J. FADILI, *Astronomical image denoising using dictionary learning*, Astronomy & Astrophysics, 556 (2013), p. A132.
- [7] Y. BENJAMINI, *Simultaneous and selective inference: Current successes and future challenges*, Biometrical Journal, 52 (2010), pp. 708–721.
- [8] Y. BENJAMINI AND Y. HOCHBERG, *Controlling the false discovery rate: A practical and powerful approach to multiple testing*, Journal of the Royal Statistical Society. Series B (Methodological), 57 (1995), pp. 289–300, <https://doi.org/10.2307/2346101>, <http://dx.doi.org/10.2307/2346101>.
- [9] S. BERNAL, F. BOTANA, J. CALVINO, C. LOPEZ-CARTES, J. PEREZ-OMIL, AND J. RODRIGUEZ-IZQUIERDO, *The interpretation of hrem images of supported metal catalysts using image simulation: profile view images*, Ultramicroscopy, 72 (1998), pp. 135–164.
- [10] A. BUADES, B. COLL, AND J.-M. MOREL, *A non-local algorithm for image denoising*, in 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), vol. 2, IEEE, 2005, pp. 60–65.
- [11] T.-O. BUCHHOLZ, M. JORDAN, G. PIGINO, AND F. JUG, *Cryo-care: content-aware image restoration for cryo-transmission electron microscopy data*, in 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), IEEE, 2019, pp. 502–506.
- [12] S. G. CHANG, B. YU, AND M. VETTERLI, *Adaptive wavelet thresholding for image denoising and compression*, IEEE Trans. Image Processing, 9 (2000), pp. 1532–1546.
- [13] Y. CHEN AND T. POCK, *Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration*, IEEE transactions on pattern analysis and machine intelligence, 39 (2016), pp. 1256–1272.

[14] P. A. CROZIER, E. L. LAWRENCE, J. L. VINCENT, AND B. D. LEVIN, *Dynamic restructuring during processing: approaches to higher temporal resolution*, Microscopy and Microanalysis, 25 (2019), pp. 1464–1465.

[15] J. M. EDE, *Deep learning in electron microscopy*, Machine Learning: Science and Technology, (2020).

[16] J. M. EDE AND R. BEANLAND, *Improving electron micrograph signal-to-noise with an atrous convolutional encoder-decoder*, Ultramicroscopy, 202 (2019), pp. 18–25.

[17] P. ERCIUS, I. JOHNSON, H. BROWN, P. PELZ, S.-L. HSU, B. DRANEY, E. FONG, A. GOLDSCHMIDT, J. JOSEPH, J. LEE, AND ET AL., *The 4d camera – a 87 kHz frame-rate detector for counted 4d-stem experiments*, Microscopy and Microanalysis, (2020), p. 1–3, <https://doi.org/10.1017/S1431927620019753>.

[18] A. FARUQI AND G. McMULLAN, *Direct imaging detectors for electron microscopy*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 878 (2018), pp. 180 – 190, <https://doi.org/10.1016/j.nima.2017.07.037>, <http://www.sciencedirect.com/science/article/pii/S0168900217307787>. Radiation Imaging Techniques and Applications.

[19] E. GIANNATOU, G. PAPAVIEROS, V. CONSTANToudis, H. PAPAGEORGIOU, AND E. GOGOLIDES, *Deep learning denoising of sem images towards noise-reduced ler measurements*, Microelectronic Engineering, 216 (2019), p. 111051.

[20] K. GONG, J. GUAN, C.-C. LIU, AND J. QI, *Pet image denoising using a deep neural network through fine tuning*, IEEE Transactions on Radiation and Plasma Medical Sciences, 3 (2018), pp. 153–161.

[21] H. GUO, P. SAUTET, AND A. N. ALEXANDROVA, *Reagent-triggered isomerization of fluxional cluster catalyst via dynamic coupling*, The Journal of Physical Chemistry Letters, 11 (2020), pp. 3089–3094, <https://doi.org/10.1021/acs.jpclett.0c00548>, <https://arxiv.org/abs/https://doi.org/10.1021/acs.jpclett.0c00548>. PMID: 32227852.

[22] Y. HEL-OR AND D. SHAKED, *A discriminative approach for wavelet denoising*, IEEE Transactions on Image Processing, 17 (2008), pp. 443–457.

[23] J. P. HORWATH, D. N. ZAKHAROV, R. MEGRET, AND E. A. STACH, *Understanding important features of deep learning models for segmentation of high-resolution transmission electron microscopy images*, npj Computational Materials, 6 (2020), pp. 1–9.

[24] S. IOFFE AND C. SZEGEDY, *Batch normalization: Accelerating deep network training by reducing internal covariate shift*, arXiv preprint arXiv:1502.03167, (2015).

[25] W. JIANG, M. L. BAKER, Q. WU, C. BAJAJ, AND W. CHIU, *Applications of a bilateral denoising filter in biological electron microscopy*, Journal of structural biology, 144 (2003), pp. 114–122.

[26] W. KHADEMI, S. RAO, C. MINNERATH, G. HAGEN, AND J. VENTURA, *Self-supervised poisson-gaussian denoising*, arXiv preprint arXiv:2002.09558, (2020).

[27] B. KIM, M. HAN, H. SHIM, AND J. BAEK, *A performance comparison of convolutional neural network-based image denoising methods: The effect of loss functions on low-dose ct images*, Medical physics, 46 (2019), pp. 3906–3923.

[28] D. P. KINGMA AND J. BA, *Adam: A method for stochastic optimization*, arXiv preprint arXiv:1412.6980, (2014).

[29] E. J. KIRKLAND ET AL., *Image simulation in transmission electron microscopy*, Cornell University, Ithaca, (2006).

[30] A. KRULL, T.-O. BUCHHOLZ, AND F. JUG, *Noise2void-learning denoising from single noisy images*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 2129–2137.

[31] A. KRULL, T. VICAR, AND F. JUG, *Probabilistic noise2void: Unsupervised content-aware denoising*, arXiv preprint arXiv:1906.00651, (2019).

[32] S. LAINE, T. KARRAS, J. LEHTINEN, AND T. AILA, *High-quality self-supervised deep image denoising*, in Advances in Neural Information Processing Systems, 2019, pp. 6970–6980.

[33] E. L. LAWRENCE, B. D. LEVIN, B. K. MILLER, AND P. A. CROZIER, *Approaches to exploring spatio-temporal surface dynamics in nanoparticles with in situ transmission electron microscopy*, Microscopy and Microanalysis, 26 (2020), p. 86–94, <https://doi.org/10.1017/S1431927619015228>.

[34] Y. LECUN, Y. BENGIO, AND G. HINTON, *Deep learning*, nature, 521 (2015), p. 436.

[35] J. LEHTINEN, J. MUNKBERG, J. HASSELGREN, S. LAINE, T. KARRAS, M. AITTALA, AND T. AILA, *Noise2noise: Learning image restoration without clean data*, arXiv preprint arXiv:1803.04189, (2018).

[36] B. D. LEVIN, E. L. LAWRENCE, AND P. A. CROZIER, *Tracking the picoscale spatial motion of atomic columns during dynamic structural change*, Ultramicroscopy, 213 (2020), p. 112978, <https://doi.org/https://doi.org/10.1016/j.ultramic.2020.112978>, <http://www.sciencedirect.com/science/article/pii/S0304399119303122>.

[37] J. W. LICHTMAN AND J.-A. CONCHELLO, *Fluorescence microscopy*, Nature methods, 2 (2005), pp. 910–919.

[38] J. S. LIM, *Two-dimensional signal and image processing*, ph, (1990).

[39] T. LINDEBERG, *Scale selection properties of generalized scale-space interest point detectors*, Journal of Mathematical

Imaging and vision, 46 (2013), pp. 177–210.

[40] F. LUISIER, T. BLU, AND M. UNSER, *A new sure approach to image denoising: Interscale orthonormal wavelet thresholding*, IEEE Transactions on Image Processing, 16 (2007), pp. 593–606.

[41] F. LUISIER, T. BLU, AND M. UNSER, *Image denoising in mixed poisson–gaussian noise*, IEEE Transactions on image processing, 20 (2010), pp. 696–708.

[42] J. MADSEN, P. LIU, J. KLING, J. B. WAGNER, T. W. HANSEN, O. WINTHER, AND J. SCHIØTZ, *A deep learning approach to identify local structures in atomic-resolution transmission electron microscopy images*, Advanced Theory and Simulations, 1 (2018), p. 1800037.

[43] M. MAKITALO AND A. FOI, *Optimal inversion of the generalized anscombe transformation for poisson-gaussian noise*, IEEE transactions on image processing, 22 (2012), pp. 91–103.

[44] B. MANIFOLD, E. THOMAS, A. T. FRANCIS, A. H. HILL, AND D. FU, *Denoising of stimulated raman scattering microscopy images via deep learning*, Biomedical optics express, 10 (2019), pp. 3860–3874.

[45] I. S. MCLEAN, *Electronic imaging in astronomy: detectors and instrumentation*, Springer Science & Business Media, 2008.

[46] W. MEINIEL, J.-C. OLIVO-MARIN, AND E. D. ANGELINI, *Denoising of microscopy images: a review of the state-of-the-art, and a new sparsity-based method*, IEEE Transactions on Image Processing, 27 (2018), pp. 3842–3856.

[47] P. MILANFAR, *A tour of modern image filtering: New insights and methods, both practical and theoretical*, IEEE signal processing magazine, 30 (2012), pp. 106–128.

[48] D. MINARIK, O. ENQVIST, AND E. TRÄGÅRDH, *Denoising of scintillation camera images using a deep convolutional neural network: a monte carlo simulation approach*, Journal of Nuclear Medicine, 61 (2020), pp. 298–303.

[49] S. MOHAN, Z. KADKHODAIE, E. P. SIMONCELLI, AND C. FERNANDEZ-GRANDA, *Robust and interpretable blind image denoising via bias-free convolutional neural networks*, in International Conference on Learning Representations, 2020, <https://openreview.net/forum?id=HJISmC4FPS>.

[50] G. MONTAVON, S. LAPUSCHKIN, A. BINDER, W. SAMEK, AND K.-R. MÜLLER, *Explaining nonlinear classification decisions with deep taylor decomposition*, Pattern Rec., 65 (2017), pp. 211–222.

[51] T. MONTINI, M. MELCHIONNA, M. MONAI, AND P. FORNASIERO, *Fundamentals and catalytic applications of ceo2-based materials*, Chemical reviews, 116 (2016), pp. 5987–6041.

[52] P. NELLIST AND S. PENNYCOOK, *Accurate structure determination from image reconstruction in adf stem*, Journal of Microscopy, 190 (1998), pp. 159–170.

[53] Y. NIE, L. LI, AND Z. WEI, *Recent advancements in pt and pt-free catalysts for oxygen reduction reaction*, Chemical Society Reviews, 44 (2015), pp. 2168–2201.

[54] J. PETERSON, J. JERNIGAN, S. KAHN, A. RASMUSSEN, E. PENG, Z. AHMAD, J. BANKERT, C. CHANG, C. CLAVER, D. GILMORE, ET AL., *Simulation of astronomical images from optical survey telescopes using a comprehensive photon monte carlo approach*, The Astrophysical Journal Supplement Series, 218 (2015), p. 14.

[55] J. PORTILLA, V. STRELA, M. J. WAINWRIGHT, AND E. P. SIMONCELLI, *Image denoising using scale mixtures of gaussians in the wavelet domain*, IEEE Trans. Image Processing, 12 (2003).

[56] M. PRAKASH, M. LALIT, P. TOMANCAK, A. KRUL, AND F. JUG, *Fully unsupervised probabilistic noise2void*, in 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI), IEEE, 2020, pp. 154–158.

[57] F. P. PREPARATA AND M. I. SHAMOS, *Convex hulls: Basic algorithms*, in Computational geometry, Springer, 1985, pp. 95–149.

[58] Y. QUAN, M. CHEN, T. PANG, AND H. JI, *Self2self with dropout: Learning self-supervised denoising from single image*, in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020, pp. 1890–1898.

[59] M. RAGONE, V. YURKIV, B. SONG, A. RAMSUBRAMANIAN, R. SHAHBAZIAN-YASSAR, AND F. MASHAYEK, *Atomic column heights detection in metallic nanoparticles using deep convolutional learning*, Computational Materials Science, 180 (2020), p. 109722.

[60] M. RAPHAN AND E. P. SIMONCELLI, *Optimal denoising in redundant representations*, IEEE Transactions on image processing, 17 (2008), pp. 1342–1352.

[61] O. RONNEBERGER, P. FISCHER, AND T. BROX, *U-net: Convolutional networks for biomedical image segmentation*, in International Conference on Medical image computing and computer-assisted intervention, Springer, 2015, pp. 234–241.

[62] D. Y. SHETH, S. MOHAN, J. L. VINCENT, R. MANZORRO, P. A. CROZIER, M. M. KHAPRA, E. P. SIMONCELLI, AND C. FERNANDEZ-GRANDA, *Unsupervised deep video denoising*, arXiv preprint arXiv:2011.15045, (2020).

[63] K. SIMONYAN, A. VEDALDI, AND A. ZISSERMAN, *Deep inside convolutional networks: Visualising image classifica-*

tion models and saliency maps, arXiv preprint arXiv:1312.6034, (2013).

[64] D. SMITH, *CHAPTER 1: Characterization of nanomaterials using transmission electron microscopy*, no. 37 in RSC Nanoscience and Nanotechnology, Royal Society of Chemistry, 37 ed., Jan. 2015, pp. 1–29, <https://doi.org/10.1039/9781782621867-00001>.

[65] G. SUN, A. N. ALEXANDROVA, AND P. SAUTET, *Structural rearrangements of subnanometer cu oxide clusters govern catalytic oxidation*, ACS Catalysis, 10 (2020), pp. 5309–5317, <https://doi.org/10.1021/acscatal.0c00824>, <https://doi.org/10.1021/acscatal.0c00824>, <https://arxiv.org/abs/https://doi.org/10.1021/acscatal.0c00824>.

[66] A. SUVEER, A. GUPTA, G. KYLBERG, AND I.-M. SINTORN, *Super-resolution reconstruction of transmission electron microscopy images using deep learning*, in 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), IEEE, 2019, pp. 548–551.

[67] F. TAO AND P. CROZIER, *Atomic-scale observations of catalyst structures under reaction conditions and during catalysis*, Chemical Reviews, 116 (2016), pp. 3487–3539, <https://doi.org/10.1021/cr5002657>.

[68] C. TOMASI AND R. MANDUCHI, *Bilateral filtering for gray and color images.*, in ICCV, vol. 98, 1998.

[69] R. K. VASUDEVAN AND S. JESSE, *Deep learning as a tool for image denoising and drift correction*, Microscopy and Microanalysis, 25 (2019), pp. 190–191.

[70] J. L. VINCENT, R. MANZORRO, S. MOHAN, B. TANG, D. Y. SHETH, E. P. SIMONCELLI, D. S. MATTESON, C. FERNANDEZ-GRANDA, AND P. A. CROZIER, *Developing and evaluating deep neural network-based denoising for nanoparticle tem images with ultra-low signal-to-noise*, (2021), <https://arxiv.org/abs/2101.07770>.

[71] Z. WANG, A. C. BOVIK, H. R. SHEIKH, AND E. P. SIMONCELLI, *Image quality assessment: from error visibility to structural similarity*, IEEE transactions on image processing, 13 (2004), pp. 600–612.

[72] Z. L. WANG, *New developments in transmission electron microscopy for nanotechnology*, Advanced Materials, 15 (2003), pp. 1497–1514.

[73] N. WIENER, *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*, Technology Press, 1950.

[74] W. YU, M. D. POROSOFF, AND J. G. CHEN, *Review of pt-based bimetallic catalysis: from model surfaces to supported catalysts*, Chemical reviews, 112 (2012), pp. 5780–5817.

[75] K. ZHANG, W. ZUO, Y. CHEN, D. MENG, AND L. ZHANG, *Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising*, IEEE Transactions on Image Processing, 26 (2017), pp. 3142–3155.

[76] X. ZHANG, Y. LU, J. LIU, AND B. DONG, *Dynamically unfolding recurrent restorer: A moving endpoint control method for image restoration*, arXiv preprint arXiv:1805.07709, (2018).

[77] Y. ZHANG, Y. ZHU, E. NICHOLS, Q. WANG, S. ZHANG, C. SMITH, AND S. HOWARD, *A poisson-gaussian denoising dataset with real fluorescence microscopy images*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 11710–11718.

[78] H. J. ZHU, B. C. HAN, AND B. QIU, *Survey of astronomical image processing methods*, in International Conference on Image and Graphics, Springer, 2015, pp. 420–429.

[79] J.-M. ZUO AND J. SPENCE, *Advanced Transmission Electron Microscopy, Imaging and Diffraction in Nanoscience*, 01 2017.

Appendix A. Data simulation.

A.1. Simulation process. The simulated TEM image dataset was generated using the multi-slice image simulation method, as implemented in the Dr. Probe software package [4]. In the multi-slice approach, the modeled specimen is sectioned into many thin slices (here they are 0.167 Angstroms thick), and quantum mechanical calculations are performed to simulate the incident electron wave function propagating through and interacting with each slice of the material [29]. The resultant wave function exiting the last slice is then convolved with a point spread function that emulates the effect of imaging it in the electron microscope. All of the image simulations were performed using an accelerating voltage of 300 kV with a beam convergence angle of 0.2 mrad and a focal spread of 4 nm. The third-order spherical aberration coefficient (C_s) was set to be $-13 \mu\text{m}$. The fifth-order spherical aberration coefficient ($C5$) was set as 5 mm. All other aberrations (e.g., 2-fold and 3-fold astigmatism, coma, star aberration, etc.) were approximated to be negligible. The defocus ($C1$) was varied systematically between 0 nm and 20 nm, as discussed below. Image calculations were computed

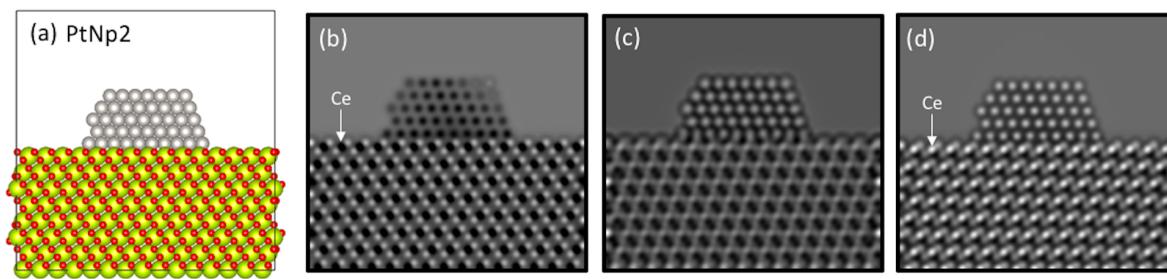


Figure A.1: Demonstration of contrast reversal with changes in defocus. (a) Image of the Pt/CeO₂ atomic structural model. (b) to (d) Simulated images under different electron-optical focusing conditions, emphasizing variations on the Ce and Pt column contrast. In (b), the image shows black contrast for both Ce and Pt columns. In (c), the Pt columns reverse contrast and now appear white, while Ce columns become challenging to discriminate. Finally, in (d) all of the atomic columns appear with white contrast.

using a non-linear model including partial temporal coherence by explicit averaging and partial spatial coherence, which is treated by a quasi-coherent approach with a dampening envelope applied to the wave function. An isotropic vibration envelope of 50 pm was applied during the image calculation. Images were simulated with 1024 x 1024 pixels and then later binned to desired sizes to match the pixel size of the experimentally acquired image series. Finally, to equate the intensity range of the simulated images with those acquired experimentally, the intensities of the simulated images were scaled by a factor which equalized the vacuum intensity in a single simulation to the average intensity measured over a large area of the vacuum in a single 0.025 second experimental frame (i.e., 0.45 counts per pixel in the vacuum region).

A.2. Experimental parameters. In phase-contrast TEM imaging (the technique employed here), multiple electron-optical and specimen parameters can give rise to complex, non-linear modulations of the image contrast. These parameters can include the objective lens defocus, the specimen thickness, the orientation of the specimen, and its crystallographic shape/structure. Due to the complex image formation mechanisms, atomic columns of the same material imaged may appear black or white (or somewhere in between, i.e., intermediate) depending on the exact combination of these various factors. Examples of the type of contrast reversal that may occur for a static structure imaged at constant thickness and tilt are given in Figure A.1. Additionally, images showing the type of contrast variations that may occur when the support thickness is changed, and how these compare to those which arise from changes in defocus are given in Figure A.2.

Here, we systematically varied these parameters to generate a large number of cases (approximately 18,000), corresponding to potential combinations that may arise during a real experiment (see Figure A.3). First, around 100 atomic-scale structural models of CeO₂-supported Pt nanoparticles were generated. Each model represents Pt nanoparticles of various size, shape and atomic structure (e.g., small, medium, or large size, with either faceted or defected surfaces, or some combination of both), supported on CeO₂, which itself may present either a faceted surface or one characterized by surface defects. Secondly, the thickness of the CeO₂ support was varied from 3 nm to 6 nm along 1 nm increments. One aspect to note is that the thickness variation is not equally applied to each of the

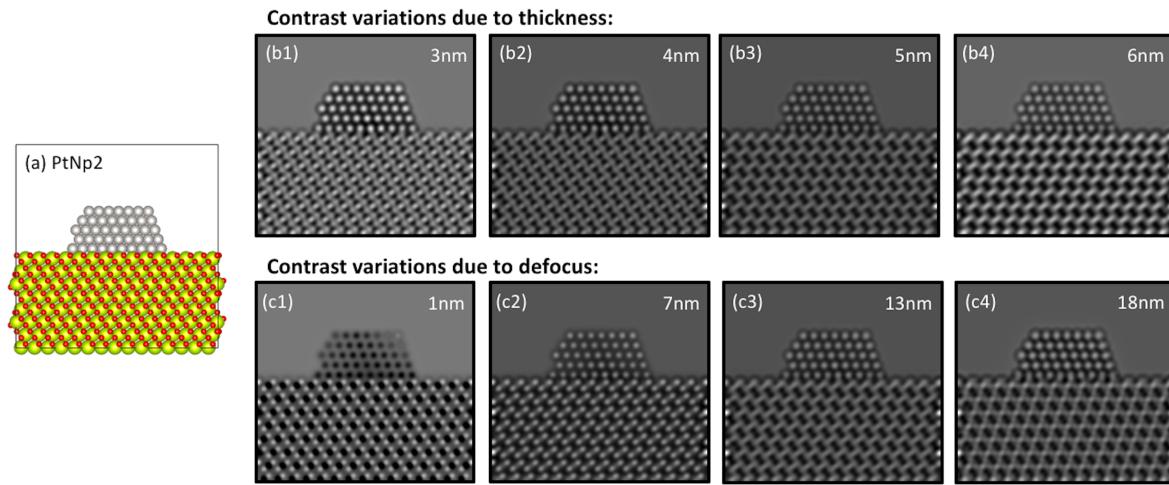


Figure A.2: Image contrast variations due to thickness and defocus. (a) Image of the Pt/CeO₂ atomic structural model. (b1) to (b4) Simulated images at a defocus value of 13 nm, where the variations of the contrast are due to the thickness of the model, increased from 3 nm to 6 nm. (c1) to (c4) Contrast variations on simulated images due to defocus: the thickness of the model has been kept constant at 5 nm and the defocus has been tuned to 1 nm, 7 nm, 13 nm, and 18 nm.

aforementioned models. Third, each resultant model was tilted from 0° to 4° about the *x* and *y* axes independently in increments of 1°. Thus, variations from 0° in *x* and 0° in *y*, to 4° in *x* and 0° in *y*, or 0° in *x* and 4° in *y* were considered. The final parameter systematically varied in the simulated image dataset was the electron optical defocus. Every model containing a unique shape/structure, thickness, and tilt (855 total) was imaged under a range of defocus values which often arise experimentally. Namely, the defocus was varied from 0 nm to 20 nm, along increments of 1 nm. Considering all combinations of the varied parameters, a total of 17,955 simulated images were generated for training and testing the neural network.

A.3. Description of nanoparticle structures. The 3D atomic structural models utilized in this work consist of faceted Pt nanoparticles that oriented in a [110] zone axis and that are supported on a CeO₂ (111) surface which is itself oriented in the [110] zone axis. This crystallographic configuration corresponds to that which is often observed experimentally and is thus the focus of the current work. All of the models have been constructed with the freely available Rhodius software [9]. Each model consists of a supercell having *x* and *y* dimensions of 5 nm x 5 nm. As discussed above, the support thickness was systematically varied for each model, and so the supercell's *z* dimension varies between 3 nm and 6 nm, depending on the thickness of the particular model.

While imaging these materials systems, experimentalists often aspire to visualize atomic-level structural rearrangements that can occur at the surfaces of the supported nanoparticles. Additionally, there are many millions of nanoparticles on a typical TEM sample, and the specific atomic-scale structural features comprising the surfaces of those imaged during an experiment may vary slightly from nanoparticle to nanoparticle. In order to encompass such complexity in the training dataset, a variety of Pt nanoparticles of multiple sizes/shape and surface defect character were incorporated into

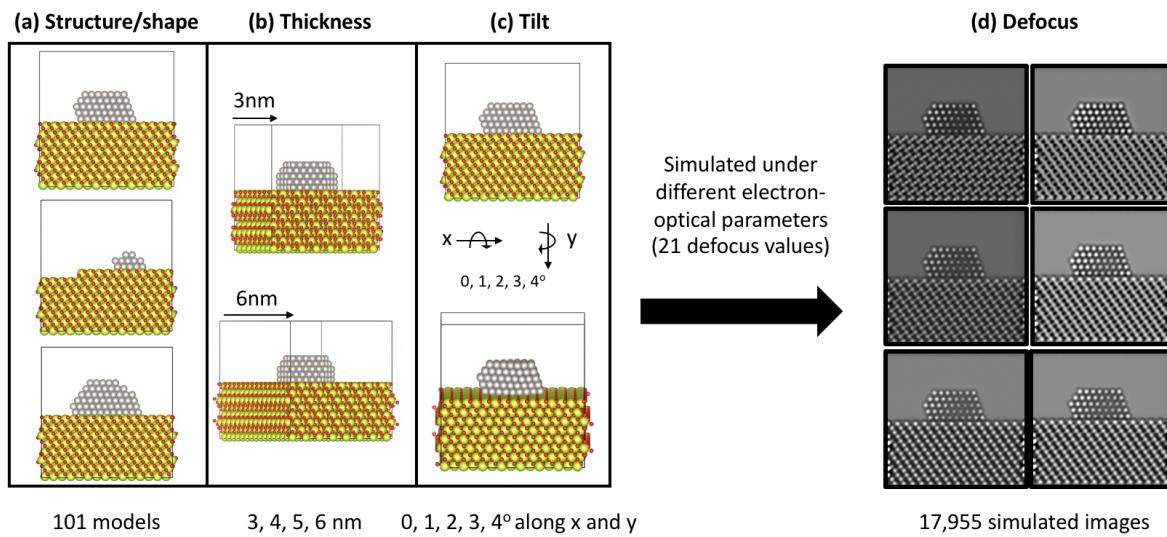


Figure A.3: Summary of parameters considered during the modelling and image simulation processes. Subset of Pt/CeO₂ atomic models presenting variations on the (a) structure and shape of the nanoparticle and the support, (b) the thickness of the CeO₂ slab and (c) the tilt of the atomic models. Color code for the models matches Pt, Ce and O with grey, yellow and red atoms respectively. (d) Simulated images under different defocus values.

the 3D models. For example, four such models of CeO₂-supported Pt nanoparticles having various size and shape are shown in parts (a) to (d) of Figure A.4. The multi-slice TEM image simulations generated from the models are shown below each for two different conditions, namely in parts (a1) to (d1), images are shown for a case in which the support is 3 nm thick, the defocus is 9 nm, and the tilt is 0° in x and 0° in y; in parts (a2) to (d2), images are shown for the case in which the support is 5 nm thick, the defocus is 6 nm, and the tilt is 4° in x and 0° in y. Furthermore, the surface character of the Pt nanoparticles was varied by altering the defect structure at different surface sites. A few examples are depicted in Figure A.5. Here, in part (a), a CeO₂-supported Pt nanoparticle with faceted surfaces is shown; directly beneath it in (a1) is an image simulated under conditions in which the support is 3 nm thick, the defocus is 9 nm, and there is no tilt. The arrowed sites designate locations on the Pt surface that have been subsequently altered. In part (b), the surface has been modified by removing a full atomic column from the arrowed location. In part (c), the occupancy of the arrowed corner site has been reduced by half. And in part (d), the occupancy of the arrowed corner site has been further reduced to a single atom. Parts (b1) - (d1) show the images simulated from these respective structures under the same imaging condition. Note that the surface sites altered in the structure correspond to high-energy sites (e.g., corners and edges) which are more likely to dynamically rearrange or show variation than, say, a low-energy terrace site located in the middle of the surface.

Appendix B. Description of denoising CNNs. In this section we describe the CNN architectures and training procedure used for our computational experiments in more detail.

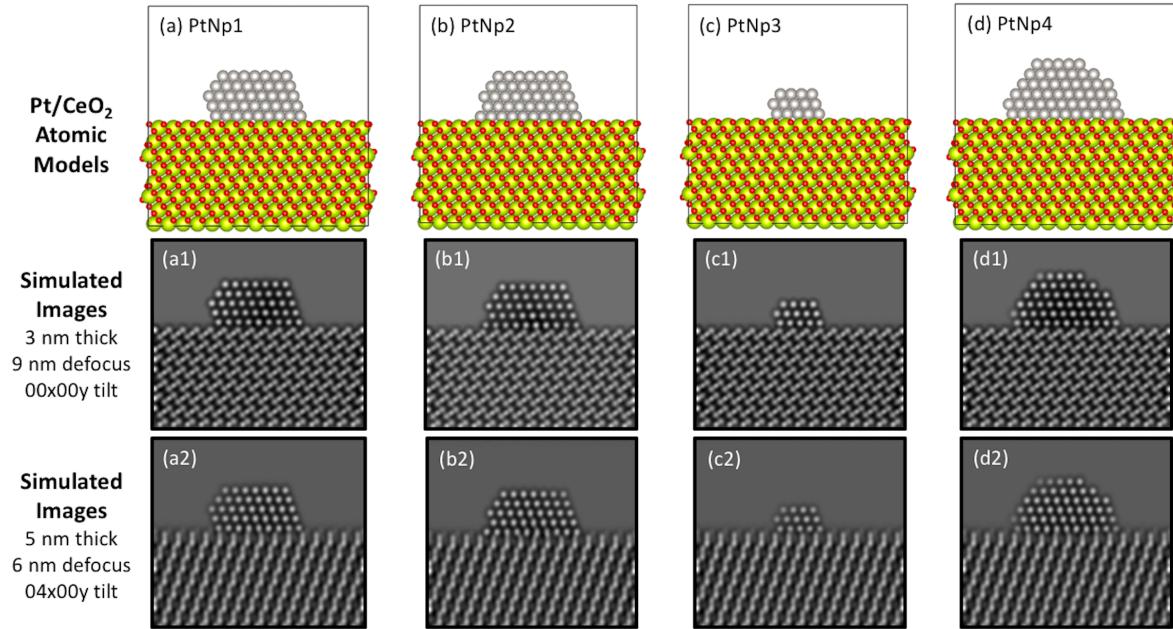


Figure A.4: **Variations in the structure/size of the supported Pt nanoparticle.** (a) to (d) Atomic models of Pt nanoparticles (grey atoms) with different shapes and supported over a CeO_2 slab (yellow and red atoms respectively). (a1) to (d1) Simulated images depicting the described atomic models, considering a thickness of 3 nm, 9 nm of defocus and no tilt on the model, whereas (a2) to (d2) illustrate the same model under different conditions: 5 nm thickness, 6 nm defocus and 4 degrees tilted along x axis. All the simulated images present a C_s value of $-13 \mu\text{m}$.

B.1. Proposed Architecture: UNet with large field of view. We propose to use a modified version of UNet [61] with $n = 4, 6$ scales to achieve a large field of view. The network consisting of n down-blocks and n up-blocks. A down-block consists of a max-pooling layer, which reduces the spatial-dimension by half, followed by a conv-block. Similarly, an up-block consists of bilinear upsampling, which enlarges the size of the feature-map by a factor of two, followed by conv-block. A conv-block consists of conv-BN-ReLU-conv-BN-ReLU, where conv represents a convolutional layer and BN stands for batch normalization [24]. In our final model, we use 128 channels in each layer of conv-block and $n = 6$ scales.

B.2. DnCNN. DnCNN [75] consists of 20 convolutional layers, each consisting of 3×3 filters and 64 channels, batch normalization [24], and a ReLU nonlinearity. It has a skip connection from the initial layer to the final layer, which has no nonlinear units.

B.3. Small UNet from DURR. We use the UNet proposed for the restoration module in DURR [76]. The architecture consists of the following:

1. *conv1* - Takes in input image and maps to 32 channels with 5×5 convolutional kernels.
2. *conv2* - Input: 32 channels. Output: 32 channels. 3×3 convolutional kernels.
3. *conv3* - Input: 32 channels. Output: 64 channels. 3×3 convolutional kernels with stride 2.

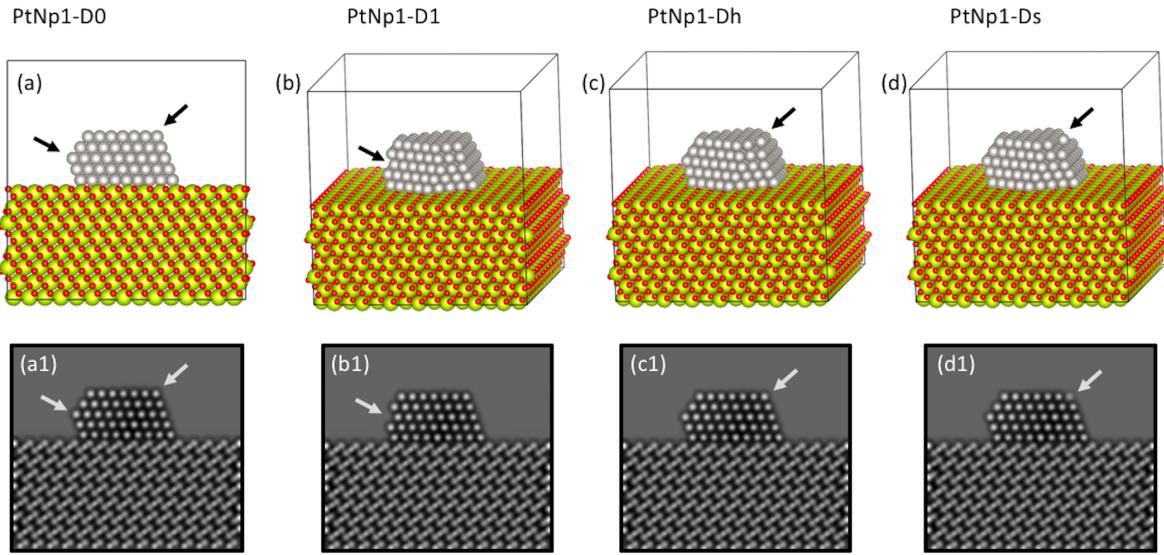


Figure A.5: Variations in the defects of the Pt surface structure. (a) Atomic model of a CeO_2 -supported Pt nanoparticle without any defects. The surface has been modified by (b) removing a full atomic column, (c) removing half of the occupancy and (d) keeping a single atom. Black arrows point the sites where these defects are taking place. Models (b), (c) and (d) have been slightly tilted to observe these modifications. (a1) to (d1) Simulated images of the presented atomic columns considering a 3nm thickness, 9 nm defocus and no tilt.

4. *conv4*- Input: 64 channels. Output: 64 channels. 3×3 convolutional kernels.
5. *conv5*- Input: 64 channels. Output: 64 channels. 3×3 convolutional kernels with dilation factor of 2.
6. *conv6*- Input: 64 channels. Output: 64 channels. 3×3 convolutional kernels with dilation factor of 4.
7. *conv7*- Transpose Convolution layer. Input: 64 channels. Output: 64 channels. 4×4 filters with stride 2.
8. *conv8*- Input: 96 channels. Output: 64 channels. 3×3 convolutional kernels. The input to this layer is the concatenation of the outputs of layer *conv7* and *conv2*.
9. *conv9*- Input: 32 channels. Output: 1 channels. 5×5 convolutional kernels.

This configuration of UNet assumes even width and height, so we remove one row or column from images in with odd height or width.

Appendix C. Additional Results. In this section we include the following additional results:

- Figure C.1 show an additional example of simulated image denoised using the proposed approach and the methods described in Sections B and 5.2.
- Figure C.2 visualizes nine nanoparticle structures used for creating the simulated dataset with surface defects described in Section 5.3.
- Figure C.3 show an additional example of real images denoised using the proposed approach and the methods described in Sections B and 5.2.

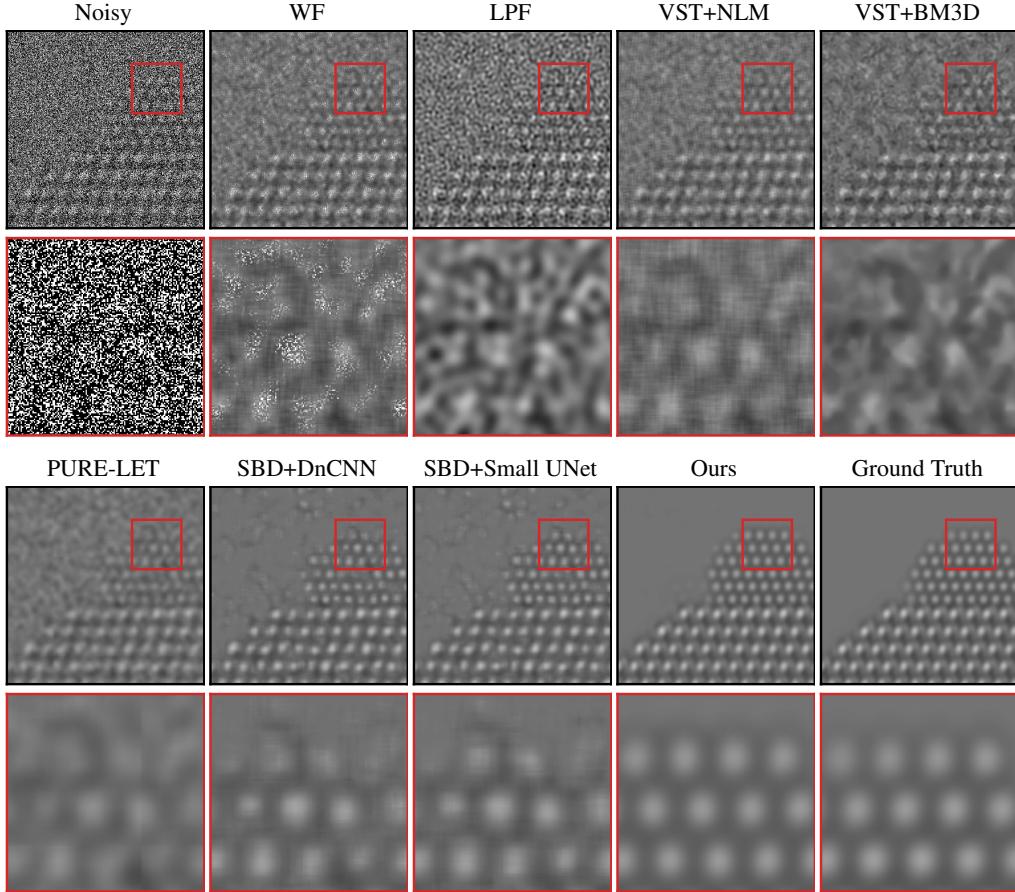


Figure C.1: **Denoising results for simulated data.** An additional example comparing SBD and the baseline methods described in Sections 5.2 and B. The second row zooms in on the region in red box. Our proposed approach produces images of much higher quality than the other approaches, and is able to accurately recover the atomic structure of the nanoparticle. For example, the vacuum region in images denoised by several of the baselines contain visible artefacts, including missing atoms.

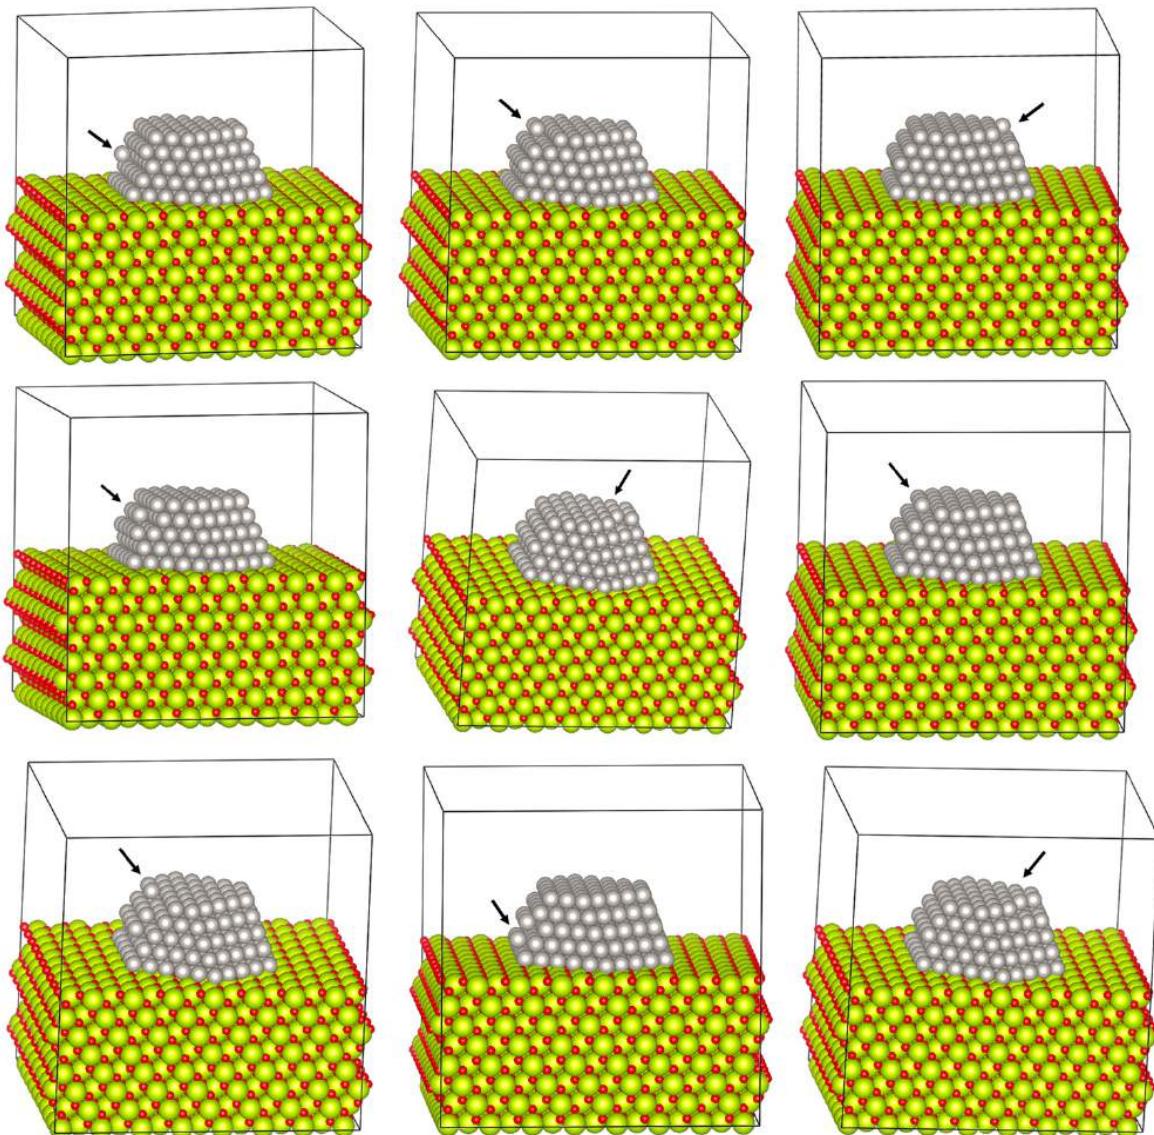


Figure C.2: Example of nanoparticle structures used for surface dataset in Section 5.3 A Subset of Pt/CeO₂ structural models with atomic-level surface defects like removal of an atom from a column, removal of two atoms, removal of all but one atom and the addition of a new atom at a site. See Section 5.3 for more details.

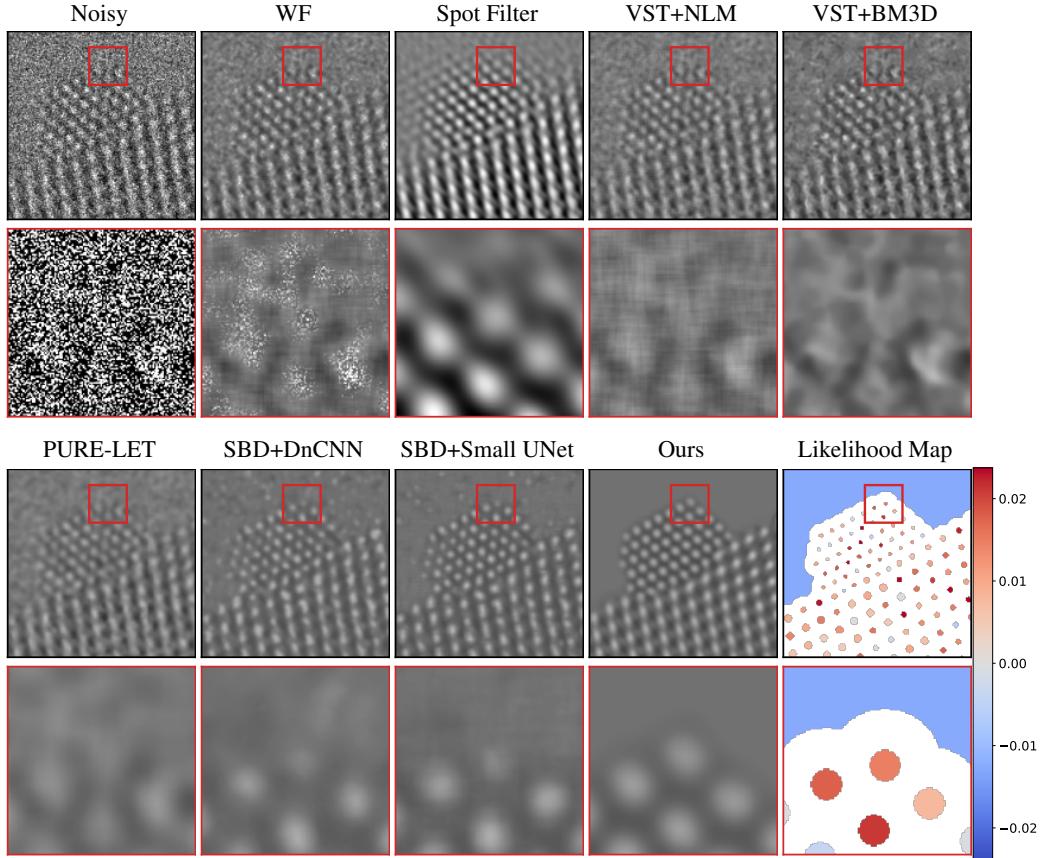


Figure C.3: Denoising results for real data. An additional example comparing SBD and the baseline methods described in Sections 5.2 and B when applied on the real data described in Section 4.1. The second row zooms in on the region in red box. In contrast to the other methods, SBD combined with the proposed architecture is able to precisely recover the structure of the nanoparticle and has very few artefacts, particularly in the vacuum region. The likelihood map quantifies the agreement between recovered structures in the denoised images, such as atomic columns and the vacuum, and the observed data (see Section 3.3 for more details).