



A mixed-scale dense convolutional neural network for image analysis

Daniël M. Pelt^a and James A. Sethian^{a,b,1}

^aCenter for Applied Mathematics for Energy Research Applications, Lawrence Berkeley National Laboratory, Berkeley, CA 94720; and ^bDepartment of Mathematics, University of California, Berkeley, CA 94720

Contributed by James A. Sethian, November 3, 2017 (sent for review September 11, 2017; reviewed by Scott E. Fraser and Ron Kimmel)

Deep convolutional neural networks have been successfully applied to many image-processing problems in recent works. Popular network architectures often add additional operations and connections to the standard architecture to enable training deeper networks. To achieve accurate results in practice, a large number of trainable parameters are often required. Here, we introduce a network architecture based on using dilated convolutions to capture features at different image scales and densely connecting all feature maps with each other. The resulting architecture is able to achieve accurate results with relatively few parameters and consists of a single set of operations, making it easier to implement, train, and apply in practice, and automatically adapts to different problems. We compare results of the proposed network architecture with popular existing architectures for several segmentation problems, showing that the proposed architecture is able to achieve accurate results with fewer parameters, with a reduced risk of overfitting the training data.

image segmentation | machine learning | convolution neural networks

Machine learning is successful in many imaging applications, such as image classification (1–3) and semantic segmentation (4–6). Many applications of machine learning to imaging problems use deep convolutional neural networks (DCNNs), in which the input image and intermediate images are convolved with learned kernels in a large number of successive layers, allowing the network to learn highly nonlinear features. The popularity of machine learning has grown significantly due to (i) recent developments that allow for effective training of deeper networks, e.g., the introduction of rectified linear units (7) and dropout layers (8); (ii) the public availability of highly optimized software to both train and apply deep networks, e.g., TensorFlow (9) and Caffe (10); and (iii) the public availability of large pretrained networks and large training datasets, e.g., VGG (2) and ImageNet (11), and will continue to be an active research area (12).

To achieve accurate results for difficult image-processing problems, DCNNs typically rely on combinations of additional operations and connections including, for example, downscaling and upscaling operations to capture features at various image scales (4, 5). To train deeper and more powerful networks, additional layer types (8, 13) and connections (14, 15) are often required. Finally, DCNNs typically use a large number of intermediate images and trainable parameters [e.g., more than 100 million (2)] to achieve results for difficult problems.

The large size and complicated nature of many DCNNs bring significant challenges. For example, the chosen combination of layers and connections can significantly influence the accuracy of trained networks. Determining which combination is best for a given problem is difficult to predict a priori. Consequently, a network that works well for one problem is not guaranteed to work well for a different problem and can require significant changes to achieve accurate results. Furthermore, the large number of parameters to learn during training requires careful choices of hyperparameters (e.g., learning rates and initialization values) to avoid problems such as overfitting (8) and vanishing gradi-

ents (13) that result in inaccurate trained networks. As a result, image analysis often relies on problem-specific traditional methods instead.

Here, we introduce a network architecture specifically designed to be easy to implement, train, and use. All layers of the network use the same set of operations and are connected to each other in the same way, removing the need to choose which operations and connections to use for each specific problem. Our proposed network architecture achieves accurate results with relatively few intermediate images and parameters, eliminating both the need to tune hyperparameters and additional layers or connections to enable training. The network uses dilated convolutions instead of scaling operations to capture features at various image scales, using multiple scales within a single layer and densely connecting all intermediate images with each other. During training, the network learns which combinations of dilations to use for the given problem, allowing the same network to be applied to different problems.

This paper is structured as follows. We first introduce notation and discuss the general structure of existing deep convolutional networks. We then introduce the proposed network architecture. We explain the experiments we performed to investigate the performance of the architecture, comparing them with popular existing architectures, and discuss their results. Finally, we conclude with a summary and final remarks.

Notation and Concepts

Problem Definition. In this paper, we apply our approach to real-valued 2D images. We define an image as a set of pixels $x \in \mathbb{R}^{m \times n \times c}$ with m rows, n columns, and c channels. We denote the image corresponding to a single channel j of x as x^j . Many

Significance

Popular neural networks for image-processing problems often contain many different operations, multiple layers of connections, and a large number of trainable parameters, often exceeding several million. They are typically tailored to specific applications, making it difficult to apply a network that is successful in one application to different applications. Here, we introduce a neural network architecture that is less complex than existing networks, is easy to train, and achieves accurate results with relatively few trainable parameters. The network automatically adapts to a specific problem, allowing the same network to be applied to a wide variety of different problems.

Author contributions: D.M.P. and J.A.S. designed research, performed research, and wrote the paper.

Reviewers: S.E.F., University of Southern California; and R.K., Technion-Israel Institute of Technology.

The authors declare no conflict of interest.

Published under the PNAS license.

¹To whom correspondence should be addressed. Email: sethian@math.berkeley.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1715832114/-DCSupplemental.




Fig. 1. A schematic representation of a two-layer CNN with input x , output y , and feature maps z_1 and z_2 . Arrows represent convolutions with nonlinear activation.

image-processing problems can be written as the problem of finding a function f that takes a certain image x and produces an output image y ; i.e., $f: \mathbb{R}^{m \times n \times c} \rightarrow \mathbb{R}^{m' \times n' \times c'}$. Note that the dimensions of the output image can be different from those of the input image. In image classification problems, for example, the output image consists of a single probability value for each of the c' possible classifications; i.e., $m' = n' = 1$. In the rest of this paper, however, we focus on problems with dense outputs, i.e., with the number of rows and columns of the output image identical to those of the input image: $m' = m$ and $n' = n$, similar to “pixel to pixel” architectures (16).

Convolutional Neural Networks. Convolutional neural networks (CNNs) model the unknown function f by using several layers that are connected to each other in succession. Each layer i produces an output image $z_i \in \mathbb{R}^{m_i \times n_i \times c_i}$, called a feature map, using output of the previous layer z_{i-1} as input. The dimensions of the layer output z_i can be different from those of the layer input z_{i-1} . The input image x is taken as the first layer z_0 , and the final layer produces the output image y .

Each individual layer can consist of multiple operations. A common layer architecture first convolves each channel of the input feature map with a different filter, then sums the resulting convolved images pixel by pixel, adds a constant value (the bias) to the resulting image, and finally applies a nonlinear operation to each pixel. These operations can be repeated using different filters and biases to produce multiple channels for the output feature map. Thus, the output z_i^j of a single channel j of such a convolutional layer is given by

$$z_i^j = \sigma(g_{ij}(z_{i-1}) + b_{ij}). \quad [1]$$

Here, $\sigma: \mathbb{R}^{m_i \times n_i} \rightarrow \mathbb{R}^{m_i \times n_i}$ is a nonlinear operation such as the popular sigmoid function or rectified linear unit (ReLU) (7), $b_{ij} \in \mathbb{R}$ is the bias, and $g_{ij}: \mathbb{R}^{m_{i-1} \times n_{i-1} \times c_{i-1}} \rightarrow \mathbb{R}^{m_i \times n_i}$ convolves each channel of the input feature map with a different filter and sums the resulting images pixel by pixel,

$$g_{ij}(z_{i-1}) = \sum_{k=0}^{c_{i-1}} C_{h_{ijk}} z_{i-1}^k, \quad [2]$$

where $C_g a$ is a 2D convolution of image a with filter g . Different ways of handling the boundaries of the image during convolution are possible: Here, we use reflective boundaries. Often, the filters h_{ijk} are relatively small (e.g., 3×3 pixels), enabling faster computation of network outputs and making the network easier to train. The architecture of the final layer can differ from other layers and can depend on the application: Common choices include using a fully connected layer instead of a convolutional one (2) or using a softmax function as the nonlinear operation for classification problems (5). A schematic of a two-layer CNN architecture is shown in Fig. 1.

The goal of training a CNN is to find filters h_{ijk} , biases b_{ij} , and potential other parameters, such that the CNN performs the task that is required. In supervised learning, training is achieved by using a set of N_t representative inputs $X = \{\hat{x}_1, \dots, \hat{x}_{N_t}\}$ with corresponding correct outputs $Y = \{\hat{y}_1, \dots, \hat{y}_{N_t}\}$ and iteratively minimizing a chosen error metric between Y and the CNN output for X . Because of the specific architecture of CNNs, partial

gradients of the error with respect to the filters and biases can be computed accurately and efficiently through backpropagation for several popular error metrics, enabling the use of efficient gradient-based optimization algorithms (17).

DCNNs. DCNNs use a network architecture similar to standard CNNs, but consist of a larger number of layers, which enables them to model more complicated functions. In addition, DCNNs often include downscaling and upscaling operations between layers, decreasing and increasing the dimensions of feature maps to capture features at different image scales. Many DCNNs incrementally downscale feature maps in the first half of the layers, called the encoder part of the network, and subsequently upscale in the second half, called the decoder part. Skip connections are often included between feature maps of the decoder and encoder at identical scales (5). A schematic representation of a common encoder–decoder DCNN architecture is shown in Fig. 2.

In general, the increased depth of DCNNs compared with shallow CNNs makes training more difficult. The increased depth often makes it more likely that training gets stuck in a local minimum of the error function and can result in gradients that become either too large or too small (13). Furthermore, DCNNs typically consist of many parameters (e.g., filters and biases), often several million or more, that have to be learned during training. The large parameter space can make training more difficult, by increasing both training time (18) and the likelihood of overfitting the network to the training data (8), thereby forcing large training sets. Several additions to standard DCNN architectures have been proposed, including batch normalization layers (13), which rescale feature maps between layers to improve the scaling of gradients during training; highway connections (14); residual connections (15); fractal networks (19), which allow information to flow more easily through deep networks by skipping layers; and dropout layers (8), in which feature maps are randomly removed from the network during training, reducing the problem of overfitting large networks.

Although these additions have advanced image processing in several fields (12), they can be difficult to routinely apply in areas such as biomedical imaging and materials science. Instead, traditional imaging algorithms are used, such as the Hough transform (20) and template matching (21), or manual processing [e.g., biological image segmentation (22)].

Theory and Algorithms

Our goal is to enable easier application of DCNNs to many imaging problems by introducing a less complicated network architecture with significantly fewer parameters to learn and which is able to automatically adapt to different problems. To do so, we introduce “the mixed-scale dense (MS-D)” network architecture, which (i) mixes scales within each layer and (ii) densely connects all feature maps.

Mixing Scales. Instead of using downscaling and upscaling operations to capture features at different scales, the MS-D




Fig. 2. A schematic representation of a common DCNN architecture with scaling operations. Downward arrows represent downscaling operations, upward arrows represent upscaling operations, and dashed arrows represent skip connections.




Fig. 3. Schematic representation of an MS-D network with $w=2$ and $d=3$. Colored lines represent 3×3 dilated convolutions, with each color representing a different dilation. Note that all feature maps are used for the final output computation.

architecture uses dilated convolutions. A dilated convolution $D_{h,s}$ with dilation $s \in \mathbb{Z}^+$ uses a dilated filter h that is nonzero only at distances that are a multiple of s pixels from the center.* Recently, it was shown that dilated convolutions are able to capture additional features in DCNNs that use the traditional scaling approach (23). Furthermore, instead of having each layer operate at a certain scale as in existing DCNNs, in the mixed-scale approach each individual channel of a feature map within a single layer operates at different scale. Specifically, we associate the convolution operations for each channel of the output image of a certain layer with a different dilation:

$$g_{ij}(\mathbf{z}_{i-1}) = \sum_{k=0}^{c_i-1} D_{h_{ijk}, s_{ij}} \mathbf{z}_{i-1}^k. \quad [3]$$

The proposed mixed-scale approach alleviates many of the disadvantages of the standard downscaling and upscaling approach. First, large-scale information about the image quickly becomes available in early layers of the network through relatively large dilations, making it possible to use this information to improve the results of deeper layers. Furthermore, information at a certain scale can be used directly to inform decisions at other scales without having to pass through layers at intermediate scales. Similar advantages were recently found when training large multi-grid architectures (24). No additional parameters have to be learned during training, since the mixed-scale approach does not include learned upscaling operations. This results in smaller networks that are easier to train. Finally, although dilations s_{ij} must be chosen in advance, the network can learn which combinations of dilations to use during training, making identical mixed-scale DCNNs applicable across different problems (experiments below).

Dense Connections. When using convolutions with reflective boundaries, the mixed-scale approach has an additional advantage compared with standard scaling: All network feature maps have the same number of rows and columns as the input and output image, i.e., $m_i = m$ and $n_i = n$ for all layers i , and hence, when computing a feature map for a specific layer, we are not restricted to using only the output of the previous layer. Instead, all previously computed feature maps $\{\mathbf{z}_0, \dots, \mathbf{z}_{i-1}\}$, including the input image \mathbf{x} , can be used to compute the layer output \mathbf{z}_i . Thus, we change the channel image computation 1 and the convolutional operation 3 to

$$\begin{aligned} \mathbf{z}_i^j &= \sigma(g_{ij}(\{\mathbf{z}_0, \dots, \mathbf{z}_{i-1}\}) + b_{ij}) \\ g_{ij}(\{\mathbf{z}_0, \dots, \mathbf{z}_{i-1}\}) &= \sum_{l=0}^{i-1} \sum_{k=0}^{c_l-1} D_{h_{ijkl}, s_{ij}} \mathbf{z}_l^k. \end{aligned} \quad [4]$$

* Alternatively, dilated convolutions can be defined without using dilated filters by changing the convolution operation itself; see ref. 23 for a detailed explanation.

Similarly, to produce the final output image \mathbf{y} , all feature maps can be used instead of only those of the last layer. We call this approach of using all previously computed feature maps densely connecting a network.

In a densely connected network, all feature maps are maximally (re)used: If a certain useful feature is detected in a feature map, it does not have to be replicated in other layers to be used deeper in the network, as in other DCNN architectures. As a result, significantly fewer feature maps and trainable parameters are required to achieve the same accuracy in densely connected networks compared with standard networks. The smaller number of maps and parameters makes it easier to train densely connected networks, reducing the risk of overfitting and enabling effective training with relatively small training sets. Recently, a similar dense-connection architecture was proposed which relied on a relatively small number of parameters (25); however, in ref. 25 the dense connections were used only within small sets of layers at a single scale, with traditional downscaling and upscaling operations to acquire information at different scales. Here, we combine dense connections with the mixed-scale approach, enabling dense connections between the feature maps of the entire network, resulting in more efficient use of all feature maps and an even larger reduction of the number of required parameters.

MS-D Neural Networks. By combining mixed-scale dilated convolutions and dense connections, we can define a DCNN architecture that we call the MS-D network architecture. Similar to existing architectures, an MS-D network consists of several layers of feature maps. Each feature map is the result of applying the same set of operations given by Eq. 4 to all previous feature maps: dilated convolutions with 3×3 pixel filters and a channel-specific dilation, summing resulting images pixel by pixel, adding a constant bias to each pixel, and finally applying a ReLU activation function. The final network output is computed with the same set of operations applied to all feature maps, using 1×1 pixel filters instead of 3×3 pixel filters. In other words, channels of the final output image are computed by taking linear combinations of all channels of all feature maps and applying an application-specific activation function to the result:

$$\mathbf{y}^k = \sigma' \left(\sum_{i,j} w_{ijk} \mathbf{z}_i^j + b'_k \right). \quad [5]$$

Different ways of choosing the number of channels per layer are possible. Here, we use a simple approach with each layer having the same number of channels, denoted by the network width w , and the number of noninput and nonoutput layers of the network denoted by the network depth d . A graphical representation of an MS-D network with $w=2$ and $d=3$ is shown in Fig. 3. The parameters that have to be learned during training are the convolution filters h_{ijkl} and biases b_{ij} of Eq. 4 and the weights




Fig. 4. (A–C) Example of the segmentation problem of the simulated dataset, with (A) the single-channel input image, (B) the correct segmentation with labels indicated by color, and (C) the output of a trained MS-D network with 200 layers.




Fig. 5. The class accuracy of a set of 100 simulated images (Fig. 4) as a function of the number of trainable parameters for the proposed MS-D network architecture and the popular U-Net architecture. For each U-Net network (U-Net- q), q indicates the number of scaling operations used. For the MS-D architecture, results are shown for dilations $s_{ij} \in [1, 10]$ (solid line) and $s_{ij} \in \{1, 2, 4, 8, 16\}$ (dashed line).

w_{ijk} and biases b'_k of Eq. 5. Given a network depth d and width w and number of input channels c_{in} and output channels c_{out} , the number of trainable parameters $N_{par} = N_{fts} + N_{wghts} + N_{bias}$ is given by $N_{fts} = 9 \sum_{i=0}^{d-1} w(iw + c_{in})$, $N_{wghts} = (wd + c_{in})c_{out}$, and $N_{bias} = wd + c_{out}$.

Compared with existing DCNN architectures, the MS-D network architecture has several advantages. Due to the mixing of scales through dilated convolutions and dense connections, MS-D networks can produce accurate results with relatively few feature maps and trainable parameters. Furthermore, an MS-D network learns which combination of dilations to use during training, allowing the same network to be effectively applied to a wide variety of problems. Finally, all layers are connected to each other in the same way and computed using the same set of standard operations, making MS-D networks easier to implement, train, and use in practice. MS-D networks do not include learned scaling operations or advanced layer types to facilitate training and do not require architecture changes when being applied to different problems. These advantages can make MS-D networks applicable beyond semantic segmentation, with potential value in classification, detection, instance segmentation, and adversarial networks (16).

Experiments

Setup. We implemented the MS-D architecture in Python, using PyCUDA (26) to enable GPU acceleration of computationally expensive parts such as convolutional operations. We note that existing frameworks such as TensorFlow (9) or Caffe (10) typically do not support the proposed mixed-scale approach well, since they assume that all channels of a certain feature map are computed in the same way. Furthermore, existing frameworks are mostly optimized for processing large numbers of relatively small images by efficiently implementing convolutions using large matrix multiplications (27). To allow the application of MS-D networks to problems with large images, we implemented the architecture using direct convolutions. Computations were performed on two workstations, with an NVidia GeForce GTX 1080 GPU and four NVidia Tesla K80 GPUs, respectively, all running CUDA 8.0.

In general, deeper networks tend to produce more accurate results than shallower networks (2). Because of the dense connections in MS-D networks, it is possible to effectively use networks that have many layers and few channels per layer, resulting in very deep networks with relatively few channels. Such very deep networks might be more difficult to train than shallower

networks, as explained above. However, we did not observe such problems and were able to use the extreme case of each layer consisting of only one channel ($w = 1$) and the number of layers d controlling the number of trainable parameters. We initialize all convolution filter parameters based on the same considerations as in ref. 3 by sampling random values from a zero-mean normal distribution with a SD of $\sqrt{2/n_c}$, where n_c is the number of incoming and outgoing connections of a feature map channel: $n_c = 9(c_{in} + w(d - 1)) + c_{out}$. All other trainable parameters are initialized to zero. Finally, in most experiments we use equally distributed dilations $s_{ij} \in [1, 10]$ by setting the dilation of channel j of layer i equal to $s_{ij} = ((iw + j) \bmod 10) + 1$.

In segmentation problems with L labels, we represent correct outputs by images with L channels, with channel j set to 1 for pixels that are assigned to label j and set to 0 for other pixels. We use the soft-max activation function in the final output layer and use the ADAM optimization method (17) during training to minimize the cross-entropy between correct outputs and network outputs (5). To compare results of MS-D networks with existing architectures for segmentation problems, we use the global accuracy metric (4), defined as the percentage of correctly labeled pixels in the network output, and the class accuracy metric (4), computed by taking the average of the true positive rates for each individual label.

Simulated Data. In a first experiment, network input consist of 512×512 -pixel single-channel images of objects with two shapes (circles and squares), 3 different sizes, and 6 possible textures, with added Gaussian noise. From all 36 combinations of shape, size, and texture, we train networks to detect 6 specific combinations, e.g., large squares with a horizontal texture, small circles with a diagonal texture, etc. We chose this segmentation problem because it requires DCNNs to combine features at small scales (pixel intensity and texture) with features at larger scales (size and shape) to produce accurate results. An example input is shown in Fig. 4A, with colors indicating the six combinations that have to be detected in Fig. 4B. We compare segmentation results of trained MS-D networks with those of the popular U-Net architecture (5): We use a TensorFlow implementation (28). U-Net architectures are similar to that shown in Fig. 2. Two main parameters influence performance: the number of downscaling (and subsequent upscaling) operations and the number of channels per feature map. We train each network with the same set of 10^5 randomly generated images, using a batch size of one image, and stop training once global accuracy for a different set of 100 validation images has not improved for 10^5 iterations.

In Fig. 5, class accuracy for an independent test set of 100 images is shown as a function of the number of trainable parameters for MS-D networks with $w = 1$ and $d \in \{25, 50, 100, 200\}$ layers and U-Net networks with two, three, four, and five scaling operations and various numbers of channels. The perfor-

Table 1. The number of trainable parameters (Pars) in millions (M), global accuracy (GA), and class accuracy (CA) for the CamVid test set

Method	Pars (M)	GA	CA
MS-D-Net (100 layers)	0.048	85.1	56.8
MS-D-Net (200 layers)	0.187	87.0	63.9
U-Net (3 scaling operations) (5)	1.863	83.2	50.4
U-Net (4 scaling operations) (5)	1.926	85.5	48.4
SegNet-Basic-EncoderAddition (4)	1.425	84.2	56.5
SegNet-Basic (4)	1.425	84.0	54.6
Boosting + Detectors + CRF (31)		83.8	62.5
Super Parsing (32)		83.3	51.2

The highest global accuracy, highest local accuracy, and smallest number of parameters out of all tested methods are shown in bold.

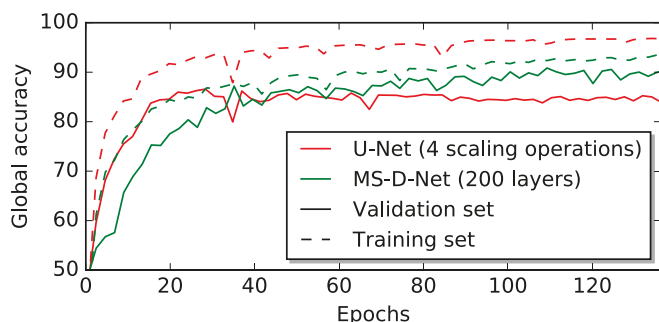


Fig. 6. The global accuracy of a U-Net network and an MS-D network as a function of the training epoch for the CamVid dataset. Given are the accuracies for the validation set (solid lines) and the training set (dashed lines).

mance of the U-Net networks depends significantly on the chosen number of scaling operations. Networks with three scaling operations are able to achieve around 80% accuracy with relatively few parameters, but do not improve significantly when using more channels per feature map, while networks with four scaling operations are able to achieve around 95% accuracy, but require a large number of parameters to do so. For a given number of parameters, MS-D networks are able to achieve significantly higher accuracies than all tested U-Net architectures, especially with relatively few parameters, and the performance of MS-D networks is similar for different choices of dilations.

CamVid Dataset. Next, we compare results for the CamVid dataset (29), using 367 training, 101 validation, and 233 testing color images of road scenes with 360×480 pixels (4). The goal is to segment 11 classes such as cars, roads, sidewalks, and pedestrians. We train MS-D networks and U-Net networks with local contrast-normalized images (30) until there is no improvement in global accuracy of the validation set, using minibatches of 10 images for MS-D networks and smaller minibatches of 3 images for U-Net networks due to memory constraints. We also report results for the SegNet architecture (4), showing the two best global accuracy results from table 1 of ref. 4, and two traditional segmentation methods (31, 32), showing the two best results from table 2 of ref. 4. For U-Net networks, the number of feature map channels was chosen such that the number of parameters was similar to that of the SegNet.

Table 1 shows global and class accuracies. MS-D segments with highest global and class accuracy, while using roughly 10 times fewer parameters. Furthermore, an MS-D network with 100 layers achieves similar accuracies to other network architectures while using 30–40 times fewer parameters.[†] Fig. 6 shows global accuracy during training for validation and training sets, for both the U-Net network and an MS-D network. Lack of improvement for the U-Net network in validation set accuracy and its difference with training set accuracy indicate overfitting of the chosen training set. Due to the smaller number of trainable parameters, the MS-D network improves validation set accuracy for more training iterations, with a significantly smaller difference with training set accuracy, showing reduced risk of overfitting of MS-D networks and the ability to accurately train with relatively small training sets. In addition, MS-D networks are able to achieve accurate results without pretraining additional large datasets, e.g., ImageNet (11), or relying on large pretrained networks, e.g., VGG (2).

[†]The authors of ref. 4 report improved results for the SegNet architecture with 90.4% global accuracy by training with a significantly larger set of around 3,500 images. However, since this larger set is not publicly available, we cannot directly compare this result with the MS-D network architecture.




Fig. 7. (A–C) A tomographic slice of the test cell (A), with the corresponding manual segmentation (B) and output of an MS-D network with 100 layers (C).

Segmenting Biomedical Images. To test whether an MS-D network can be easily applied to a new problem without adjustments, we use the same network parameters as above, with $w = 1$, $d = 100$, and dilations $s_{ij} \in [1, 10]$ applied to segmenting cell structures. We use eight manual segmentations of $512 \times 512 \times 512$ tomographic reconstructions of (mouse) lymphoblastoid cells, consisting of five labels: nuclear envelope, euchromatin, heterochromatin, mitochondria, and lipid drops. A sample tomographic slice and corresponding manual segmentation are shown in Fig. 7 A and B. The labeling of cell structures depends on multiple factors at different image scales, such as the position of the structure relative to other structures, and the pixel intensity differences between two structures can be relatively small, making it difficult to use traditional methods to perform automatic labeling. Instead, researchers rely on time-consuming manual segmentation.

To learn limited 3D features, we use five channels in the input image of the MS-D network: the current slice to be segmented and four adjacent slices. Of eight manual cell segmentations, we randomly chose six for training and one for validation and report results for the remaining cell. During training, we used a batch size of 10 images and stopped after no improvements in global accuracy for the validation cell, yielding network parameters with the best global accuracy. Fig. 7C shows network output for the slice in Fig. 7A, showing high similarity to manual segmentation. The remaining differences between network output and manual segmentation, indicated by an arrow in Fig. 7, typically represent ambiguous cell structure (see Figs. S1 and S2 for additional results). Final global accuracy and class accuracy of the trained network for the test cell are 94.1% and 93.1%, indicating that identical MS-D networks can be trained for different problems. Results for two other challenging problems are given in Figs. S3 and S4.




Fig. 8. (A–C) Tomographic images of a fiber-reinforced minicomposite, reconstructed using 1,024 projections (A) and 128 projections (B). In C, the output of an MS-D network with image B as input is shown. Bottom Right Insets in A–C show enlarged images of small regions indicated by red squares.

Denosing Large Tomographic Images. Finally, we use the above architecture, changing only the nonlinear function of the final layer from the soft-max function to the identity, and train on the different task of denoising tomographic reconstructions of a fiber-reinforced minicomposite. A total of 2,160 images of $2,560^2$ pixels were reconstructed using 1,024 acquired X-ray projections to obtain images with relatively low amounts of noise (Fig. 8A). Noisy images of the same object were obtained by reconstructing using 128 projections (Fig. 8B). The input is a noisy image, with a corresponding noiseless image used as target output during training. From the sample top, 500 images were used for training, and 100 images were used for validation. Fig. 8C shows output for a tested image near the sample bottom, computed in 2.05 s using a GTX 1080 GPU (see Fig. S5 for additional timings). The MS-D network accurately denoises highly noisy images by learning image features from the training set, and identical MS-D networks can be easily applied to different problems with minimal changes.

Conclusions

We have presented a deep convolutional MS-D network architecture for image-processing problems, using dilated convolutions instead of traditional scaling operations to learn features at different scales, using multiple scales in each layer, and computing

the feature map of each layer using all feature maps of earlier layers, resulting in a densely connected network. By combining dilated convolutions and dense connections, the MS-D network architecture can achieve accurate results with significantly fewer feature maps and trainable parameters than existing architectures, enabling accurate training with relatively small training sets. MS-D networks are able to automatically adapt by learning which combination of dilations to use, allowing identical MS-D networks to be applied to a wide range of different problems.

ACKNOWLEDGMENTS. Supporting contributions were performed by A. Ekman, C. Larabell [supported by the National Institutes of Health-National Institute for Drug Abuse (NIH-NIDA) Grant U01DA040582], S. Mo, O. Jain, D. Parkinson, A. MacDowell, and D. Ushizima [Center for Advanced Mathematics for Energy Research Applications (CAMERA)]. Computer calculations were performed at the Lawrence Berkeley National Laboratory under Contract DE0AC02-5CH11231. Tomographic images of fiber-reinforced minicomposite were provided by N. Larson and collected at the US Department of Energy's (DOE) Advanced Light Source (ALS) Beamline 8.3.2. This work was supported by CAMERA, jointly funded by The Office of Advanced Scientific Research (ASCR) and the Office of Basic Energy Sciences (BES) within the DOE's Office of Science. Soft X-ray tomography data were collected and segmented at the National Center for X-ray Tomography, supported by the National Institutes of Health-National Institute of General Medical Sciences (NIH-NIGMS) Grant P41GM103445 and the Department of Energy, Office of Biological and Environmental Research, Grant DE0AC02-5CH11231.

- Agrawal P, Girschick R, Malik J (2014) Analyzing the performance of multilayer neural networks for object recognition. *European Conference on Computer Vision 2014, Part VII, Lecture Notes in Computer Science*, eds Fleet D, Pajdla T, Schiele B, Tuytelaars T (Springer, Cham, Switzerland), pp 329–344.
- Simonyan K, Zisserman A (2014) Very deep convolutional networks for large-scale image recognition. arXiv:1409.1556.
- He K, Zhang X, Ren S, Sun J (2015) Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *Proceedings of the IEEE International Conference on Computer Vision* (Institute of Electrical and Electronics Engineers, New York), pp 1026–1034.
- Badrinarayanan V, Kendall A, Cipolla R (2016) SegNet: A deep convolutional encoder-decoder architecture for image segmentation. arXiv:1511.00561v3.
- Ronneberger O, Fischer P, Brox T (2015) U-net: Convolutional networks for biomedical image segmentation. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, eds Navab N, Hornegger J, Wells W, Frangi A (Springer, Cham, Switzerland), pp 234–241.
- Shelhamer E, Long J, Darrell T (2017) Fully convolutional networks for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39: 640–651.
- Nair V, Hinton GE (2010) Rectified linear units improve restricted Boltzmann machines. *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, eds Fürnkranz J, Joachims T (OmniPress, Madison, WI), pp 807–814.
- Srivastava N, Hinton GE, Krizhevsky A, Sutskever I, Salakhutdinov R (2014) Dropout: A simple way to prevent neural networks from overfitting. *J Mach Learn Res* 15: 1929–1958.
- Abadi M, et al. (2016) Tensorflow: A system for large-scale machine learning. *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, eds Keeton K, Roscoe T (USENIX Association, Berkeley, CA), pp 265–283.
- Jia Y, et al. (2014) Caffe: Convolutional architecture for fast feature embedding. arXiv:1408.5093.
- Deng J, et al. (2009) Imagenet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition, 2009 (CVPR 2009)* (Institute of Electrical and Electronics Engineers, New York), pp 248–255.
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521:436–444.
- Ioffe S, Szegedy C (2015) Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv:1502.03167.
- Srivastava RK, Greff K, Schmidhuber J (2015) Training very deep networks. *Adv Neural Inf Process Syst* 2377–2385.
- He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York), pp 770–778.
- Isola P, Zhu JY, Zhou T, Efros AA (2016) Image-to-image translation with conditional adversarial networks. arXiv:1611.07004.
- Kingma D, Ba J (2014) Adam: A method for stochastic optimization. arXiv:1412.6980.
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, eds Pereira F, Burges CJC, Bottou L, Weinberger KQ (Curran Associates, Inc., Red Hook, NY), Vol 25, pp 1097–1105.
- Larsson G, Maire M, Shakhnarovich G (2016) Fractalnet: Ultra-deep neural networks without residuals. arXiv:1605.07648.
- Ballard DH (1981) Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognit* 13:111–122.
- Lee Y, Hara T, Fujita H, Itoh S, Ishigaki T (2001) Automated detection of pulmonary nodules in helical CT images based on an improved template-matching technique. *IEEE Trans Med Imaging* 20:595–604.
- Gerig G, Jomier M, Chakos M (2001) Valmet: A new validation tool for assessing and improving 3D object segmentation. *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2001*, eds Niessen WJ, Viergever MA (Springer, London), pp 516–523.
- Yu F, Koltun V (2016) Multi-scale context aggregation by dilated convolutions. arXiv:1511.07122.
- Ke TW, Maire M, Yu SX (2017) Multigrid neural architectures. *The IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York), pp 4067–4075.
- Huang G, Liu Z, van der Maaten L, Weinberger KQ (2017) Densely connected convolutional networks. *The IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York), pp 2261–2269.
- Klöckner A, et al. (2012) PyCUDA and PyOpenCL: A scripting-based approach to GPU run-time code generation. *Parallel Comput* 38:157–174.
- Chetlur S, et al. (2014) cudnn: Efficient primitives for deep learning. arXiv:1410.0759.
- Akeret J, Chang C, Lucchi A, Refregier A (2017) Radio frequency interference mitigation using deep convolutional neural networks. *Astron Comput* 18: 35–39.
- Brostow GJ, Fauqueur J, Cipolla R (2009) Semantic object classes in video: A high-definition ground truth database. *Pattern Recognit Lett* 30:88–97.
- Jarrett K, et al. (2009) What is the best multi-stage architecture for object recognition? 2009 IEEE 12th International Conference on Computer Vision (Institute of Electrical and Electronics Engineers, New York), pp 2146–2153.
- Ladický L, Sturges P, Alahari K, Russell C, Torr P (2010) What, where and how many? Combining object detectors and CRFs. *Comput Vis—ECCV* 2010 424–437.
- Tighe J, Lazebnik S (2013) Superparsing. *Int J Comput Vis* 101:329–349.