

# Towards Better Analysis of Deep Convolutional Neural Networks

Mengchen Liu, Jiaxin Shi, Zhen Li, Chongxuan Li, Jun Zhu, Shixia Liu

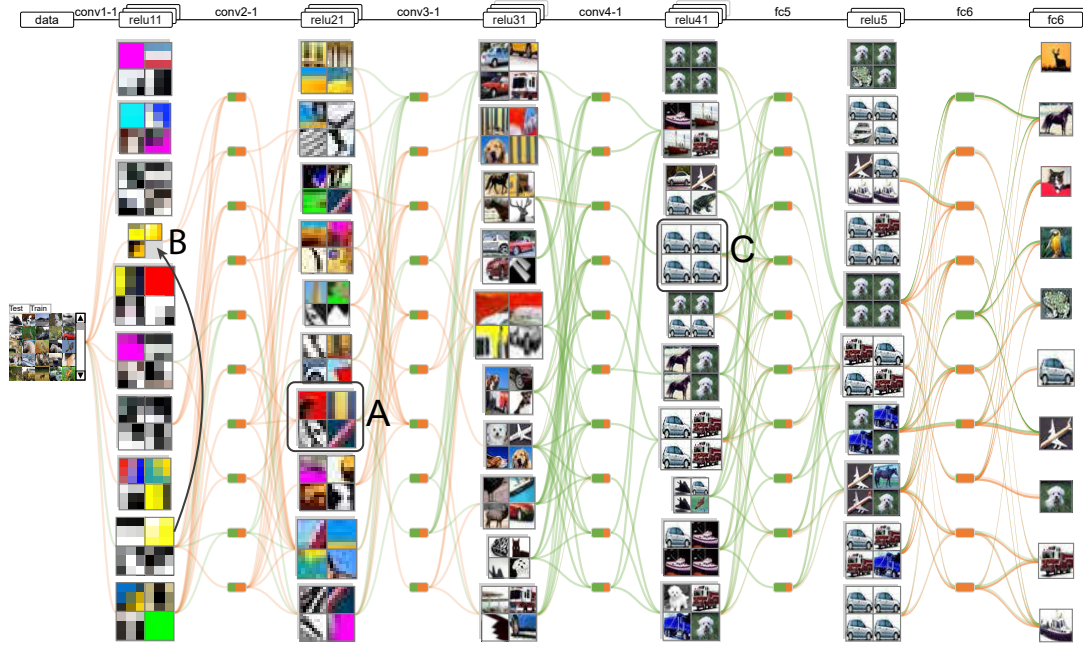


Fig. 1. CNNVis, a visual analytics toolkit that helps experts understand, diagnose, and refine deep CNNs.

**Abstract**— Deep convolutional neural networks (CNNs) have achieved breakthrough performance in many pattern recognition tasks such as image classification. However, the development of high-quality deep models typically relies on a substantial amount of trial-and-error, as there is still no clear understanding of when and why a deep model works. In this paper, we present a visual analytics approach for better understanding, diagnosing, and refining deep CNNs. We formulate a deep CNN as a directed acyclic graph. Based on this formulation, a hybrid visualization is developed to disclose the multiple facets of each neuron and the interactions between them. In particular, we introduce a hierarchical rectangle packing algorithm and a matrix reordering algorithm to show the derived features of a neuron cluster. We also propose a biclustering-based edge bundling method to reduce visual clutter caused by a large number of connections between neurons. We evaluated our method on a set of CNNs and the results are generally favorable.

**Index Terms**—deep convolutional neural networks, rectangle packing, matrix reordering, edge bundling, biclustering.

## 1 INTRODUCTION

Deep convolutional neural networks (CNNs) have demonstrated significant improvements over traditional approaches in many pattern recognition tasks [31], such as speech recognition [42, 43], image classification [17, 30], and video classification [27, 58]. More recently, deep CNNs have been employed as function approximators in deep reinforcement learning to extract robust representations and help make decisions, which has led to human-level performance in intelligent tasks such as Atari games [37] and the game of Go [44]. However, a deep CNN is often treated as a “black box” model because of its incomprehensible functions and unclear working mechanism [4]. It is generally difficult for machine learning experts to understand the role of each component (neuron, connection) due to the large number of interacting, non-linear parts in a CNN. Without a clear understanding of how and why these networks work, the development of high-performance mod-

els typically relies on a substantial amount of trial-and-error [3, 4, 57], which is time-consuming. For example, training a single deep CNN on a large dataset may take several days or even weeks.

There are two technical challenges to understanding and analyzing deep CNNs. First, a CNN may consist of tens or hundreds of layers (depth), thousands of neurons (width) in each layer, as well as millions of connections between neurons. Such large CNNs are hard to study due to the sizes involved. Second, CNNs consist of many functional components whose values and roles are not well understood either as individuals or as a whole [24]. In addition, how the non-linear components interact with each other and with other linear components in a CNN is not well understood by experts. In most cases, it is hard to summarize reusable knowledge from a failed or successful training case and transfer it to the development of other relevant deep learning models.

To tackle these challenges, we have developed an interactive, visual analytics system called CNNVis, which aims to help machine learning experts better understand, diagnose, and refine CNNs. Based on the characteristics of a deep CNN, we formulate it as a directed acyclic graph (DAG), in which each node represents a neuron and each edge represents the connection between a pair of neurons. In order to visualize a large CNN, we first cluster the layers in the network and select a representative one from each layer cluster. Then we cluster neurons in each representative layer and select several representative neurons

• M. Liu, J. Shi, Z. Li, C. Li, J. Zhu, and S. Liu are with Tsinghua University.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of

Publication xx xxx. 201x; date of current version xx xxx. 201x.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.

Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx/

from each neuron cluster. On the basis of the DAG representation, we develop a hybrid visualization to disclose the interactions between neurons and the multiple facets of each neuron by indicating its role for different types of images. In particular, a hierarchical rectangle packing algorithm is developed to show the derived features of the neuron cluster. We also design a matrix reordering algorithm based on the Held-Karp algorithm (state compression dynamic programming) [18] to demonstrate the cluster patterns in the activations of each neuron cluster. Here, the activation is the output value of a neuron, which is determined by the activation function that transforms the input value to the output value of the neuron. Moreover, we propose a biclustering-based edge bundling method to reduce the visual clutter caused by the large number of connections between neurons.

In this work, we use image classification as an example and conduct three case studies with machine learning experts. In particular, the first case study helps to illustrate the influence of the CNN model structure on performance, especially the depth and width of a CNN; the second case study demonstrates how CNNVis helps diagnose the potential issues of a failed training case; and the last case study illustrates how CNNVis helps refine a CNN to improve its performance. The case studies have shown that with CNNVis, experts can better explore and understand a deep CNN, including the role of each neuron and the connections between neurons. For example, if a CNN suffers from overfitting in the training process, some neurons learn the same feature(s), which indicates that some of them are redundant. Furthermore, experts can diagnose the potential issues of the model structure and refine a CNN, which enables more rapid iteration and faster convergence in model construction.

The key technical contributions of this work are:

- **A visual analytics system** that helps experts understand, diagnose, and refine deep CNNs.
- **A hybrid visualization** that combines a DAG with rectangle packing, matrix visualization, and a biclustering-based edge bundling method.

## 2 RELATED WORK

To help experts better understand a deep CNN, researchers in the field of computer vision have made efforts to illustrate the learned features of each neuron, which is represented by part of a real image or a synthesized image. Existing methods can be classified into two categories, namely, code inversion [12, 35, 58] and activation maximization [13, 41, 45, 53, 57].

Code inversion methods synthesize an image from the activation vector of a specific layer, which is produced by a real image. For example, Zeiler et al. [58] utilized a multi-layered Deconvolutional Network [59] to project the activations onto the input pixel space. However, simple projection without considering any prior will produce images that do not resemble natural images. To solve this problem, Mahendran et al. [35] proposed incorporating several natural image priors like  $\alpha$ -norm and total variation to make the reconstructed images more realistic. Recently, Dosovitskiy et al. [12] trained a CNN to reconstruct the images from the activations. They argued that a CNN can learn more powerful priors and have better performance than that of the manually defined ones.

Activation maximization methods aim to find an image that maximally activates a given neuron. It can be modeled as an optimization problem over the image space. Similar to code inversion methods, natural image priors are necessary as regularization during the optimization to obtain realistic images. As a result, most activation maximization methods focus on defining the regularization term using natural image priors [13, 41, 53, 57]. For example, Erhan et al. [13] constrained the L2-norm of the image to be constant. Yosinski et al. [57] defined several more powerful priors, including Gaussian blur, clipping pixels with a small norm, and clipping pixels with a small contribution.

The aforementioned methods employ a grid-based representation to display the neuron features. Although they can show the reconstructed intermediate states of each layer, they fail to disclose the inner working mechanism of CNNs, especially the role of each neuron for different types of images and the interactions between neurons. Unlike these methods, we formulate a deep CNN as a DAG. Based on the DAG

representation, we have developed a hybrid visualization that consists of rectangle packing, matrix ordering, and biclustering-based edge bundling. Empowered by the hybrid visualization, our visual analytics approach well discloses the multiple facets of each neuron as well as the interactions between them, which is very useful to understand the inner working mechanism of a deep CNN.

In the field of visualization, researchers have achieved a great deal of success in modeling domain-specific data as a DAG. Typical data includes dynamic relationships between entities [34, 50, 51], temporal topic data [10, 14, 47, 56], temporal event sequences [54], evolving egocentric network [55], and the information of musicians [23]. Researchers have also developed a set of visualizations to reveal patterns learned from the above data. However, none of these visualizations can be directly applied to illustrate deep CNNs because they lack a way to efficiently handle a large CNN that consists of tens or hundreds of layers, thousands of neurons in each layer, and millions of connections between neurons. In addition, these methods do not disclose the multiple facets of each neuron by showing its role for different types of images.

Another relevant method is BiSet [48], which employs biclustering-based edge bundling to explore coordinated relationships between entity sets. In BiSet, each edge is unweighted; while in a deep CNN, each edge has a weight to indicate the impact of the input on the output. If we simply convert a CNN to an unweighted graph and then use the biclustering method in BiSet, we may lose some important biclusters. To solve this problem, we have developed a weighted biclustering method based on the Apriori algorithm, which is an algorithm for frequent item set mining [2].

Similar to our work, Tzeng et al. [52] also employed a DAG to represent a neural network. Although this visualization method can illustrate the interactions between neurons, it suffers from serious visual clutter when handling large neural networks. To address this issue, we first cluster the layers in the network and select a representative from each layer cluster. Then we cluster neurons in each representative layer and select several representative neurons from each neuron cluster. Each node in the DAG represents a neuron cluster and the edge between nodes represents the connection between the neurons in each cluster. We have also proposed a biclustering-based edge bundling method to reduce visual clutter caused by a large number of connections between neurons.

## 3 BACKGROUND




Fig. 2. The typical architecture of a CNN.

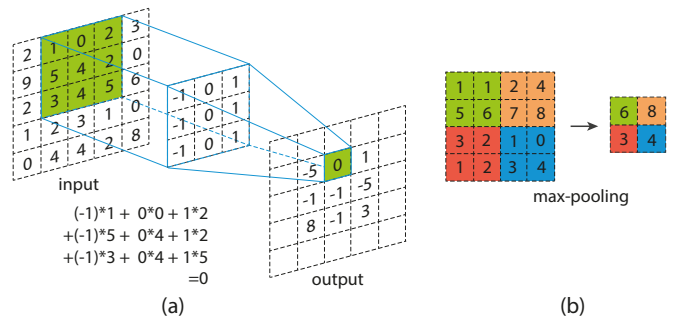


Fig. 3. Illustration of convolution and max-pooling: (a) convolution; (b) max-pooling.

In this section, we briefly introduce the architecture of CNNs and several basic concepts, which are useful for subsequent discussions.

CNNs are a specialized kind of neural networks for processing data that have a known, grid-like topology [31]. First, we briefly illustrate the architecture of CNNs.

**Architecture.** As shown in Fig. 2, a CNN is typically composed of multiple alternating convolutional and pooling layers, followed by one or several fully connected layers [31]. CNNs exploit local correlations by enforcing a local connectivity pattern between neurons of adjacent layers, namely, the inputs of neurons in the current layer come from a subset of neurons in the previous layer. This hierarchical architecture allows convolutional neural networks to extract more and more abstract representations from the lower layer to the higher layer. Fig. 2 illustrates the architecture of a CNN that contains two convolutional and two pooling layers followed by one fully connected layer. Next, we introduce the key components of CNNs.

**Convolution.** A convolution operation is performed as a window of weights slides across an image, where an output pixel produced at each position is a weighted sum of the input pixels covered by the window. The weights that parameterize the window remain the same throughout the scanning process. Therefore, convolutional layers can capture the shift-invariance of visual patterns and learn robust features. The convolution operation is illustrated in Fig. 3(a), where the value of the green pixel in the output is the weighted sum of the pixels in the green region of the input.

**Activation Function.** An activation function is a non-linear transformation that has been traditionally used in neural networks. For convolutional layers, the activation function is applied after the convolution operation. By employing activation functions, CNNs avoid learning trivial linear combinations of the inputs. One of the most popular activation functions is the rectified linear unit (ReLU) [39]. This activation function is a piecewise linear function that prunes the negative part of the input to zero and retains the positive part:

$$f(x) = \max(0, x). \quad (1)$$

For classification tasks using probability-based loss functions like cross-entropy (see the loss function part), we often require the network output to be a vector of label probabilities, which add up to 1. The softmax function is a special kind of activation function that satisfies this constraint:

$$f(x)_i = \frac{e^{x_i}}{\sum_j e^{x_j}}, \quad (2)$$

where  $x$  is the result of the linear transformation through the weights in the output layer. After applying softmax, the output  $f(x)$  is normalized to add up to 1.

**Pooling.** A pooling operation computes a specific norm over small regions on the input, which achieves some level of translation invariance. This operation aggregates small patches of pixels and thus downsamples the image features from the previous layer, which significantly reduces the computational cost when the neural network is deep. The most commonly used pooling operation in CNNs is max-pooling, which outputs the maximum ( $L_\infty$  norm) pixel value of the input region (Fig. 3(b)).

**Normalization.** Normalization is an optional operation in CNNs. It is used to speed up the convergence of the training process and reduce the probability of getting stuck in local optima [22]. This operation works by normalizing the output of certain layers through linear or non-linear operations. Many normalization methods have been developed for CNNs such as batch normalization [22].

**Loss Function.** A loss function is used to evaluate the difference between the output of a CNN and a true image label (i.e., the loss). The aim of training a CNN is to minimize the value of the loss function. It is usually achieved with stochastic gradient descent [6], an optimization method that calculates the gradient of the loss function with respect to the weight of each edge in the network and then updates the weight according to the computed gradient. Among various kinds of loss functions, the cross-entropy loss along with softmax output activations is most commonly used for classification tasks. This loss calculates the cross entropy between the ground truth distribution and the predicted distribution of CNNs:

$$l_c(o, t) = \sum_{i=1}^n -t_i \log o_i, \quad (3)$$

where  $n$  is the number of classes,  $o$  denotes the network output that is represented by a vector of probabilities for each class label, and  $t$  is a one-hot vector for the true label of the current input.

Another commonly used loss function is the hinge loss, which measures the difference between the score of the correct class and the score of the predicted class. It is likely to have better performance on object detection tasks [15]. For the sake of simplicity, we introduce the hinge loss function for two classes, which is defined by:

$$l_h(o, t) = \max(0, 1 - t \cdot o), \quad (4)$$

where  $t \in \{-1, +1\}$  is the class label, and  $o$  is the real-valued class score produced by the network. The extension to a multi-class hinge loss can be found in [32].

## 4 CNNVis

CNNVis was designed with a team of deep learning experts (six researchers) over the course of twelve months. For simplicity's sake, we denote these experts as  $E_i$  ( $i = 1, 2, \dots, 6$ ). We held discussions every two weeks. Three co-authors of this paper are also members of the team. The development of CNNVis was triggered by their need to make sense of the inner mechanisms of deep CNNs and their dissatisfaction with the state-of-the-art toolkits.

Common deep learning frameworks include Caffe [25], Theano [5], Torch [7], and TensorFlow [1]. Researchers can use these frameworks to train, debug, and deploy CNNs. Although the deep learning frameworks output high-level statistical information, such as the training loss, as well as debugging information, such as the learned features of neurons and the gradients of weights, it fails to disclose the role of each neuron for different categories of images and how the neurons work together. Accordingly, if a training process fails, it is still hard for experts to figure out what is wrong with the current model design. The experts have expressed that the development of high-quality CNN models is usually a trial-and-error procedure. As a result, they need a toolkit that can help them better understand the inner mechanism of CNNs, including the role of each neuron for the different categories of images as well as the interactions between neurons. This will allow them to summarize reusable knowledge from a failed or successful training case and transfer it to other relevant deep learning tasks.

### 4.1 Requirement Analysis

We identified the following high-level requirements based on our discussions with the experts and previous research.

**R1 - Providing an overview of the learned features of neurons.** All the experts commented that an overview of the learned features of neurons is necessary to begin their analysis (e.g., diagnosis or refinement of the model). They usually examine the quality of each learned feature layer by layer to discover potential problems. However, such an examination can be very difficult for a deep CNN with tens or hundreds of layers and thousands of neurons in a layer. As a result, they stated the need to cluster neurons into clusters so they can gain a quick overview of the learned features of each cluster.

**R2 - Interactively modifying the neuron clustering results.** Since the clustering algorithm may be imperfect and different users may have different needs, experts need to interactively modify the clustering results based on their knowledge. Expert  $E_2$  commented that when examining the training results of a CNN, he found a neuron for detecting a color patch in a cluster that mainly consists of neurons for detecting stripes with various orientations. To increase the clustering accuracy and better compare these clusters, he moved the neuron to a cluster that mainly consists of neurons for detecting color patches.

**R3 - Exploring multiple facets of neurons.** Previous work mainly focused on visualizing the learned features of neurons. In addition to this feature, the experts also requested viewing other facets of neurons. For example, expert  $E_1$  said, "In addition to the learned features, other numerical features such as activation (of a neuron) can also help me better understand its role in a classification task." During the discussion, we gradually identified that the major facets of interest are the learned features (all the experts), activations ( $E_1, E_3, E_4, E_5, E_6$ ), and contributions




Fig. 4. CNNVis pipeline.

to the final result (all the experts). Visually illustrating them can help experts gain a more comprehensive understanding of the roles of neurons.

**R4 - Revealing how low-level features are aggregated into high-level features.** In a CNN, neurons in lower layers learn to detect simple features such as stripes or corners, neurons in middle layers learn to detect a part of an object, and neurons in higher layers learn to detect a concept (e.g., a cat). This is achieved with a local connectivity pattern between neurons of adjacent layers, which means the inputs of neurons in layer  $m$  are from a subset of neurons in layer  $m-1$ . As a result, the experts wanted to learn how neurons in adjacent layers interact with each other and aggregate the low-level features into high-level features. Previous research has also shown that analyzing such connections can help experts understand how a large number of non-linear parts interact with each other [52]. A large CNN may contain millions of connections between neurons. If we display all them, it is difficult to discern individual connection due to visual clutter caused by excessive edges and edge crossings. Thus, the experts required to examine the major trends among these connections.

**R5 - Examining the debugging information.** In the discussions, the experts expressed the need to examine the debugging information of the deep model. Expert  $E_3$  said, “I often examine the debugging information such as the gradients, to diagnose a training process that failed to converge.” In addition to gradients, showing other derived values such as the relative change of weights, has also been requested by the experts. The debugging information is usually huge. For example, there are millions of gradients. It is very hard to examine them one by one and develop a full understanding. As a result, the experts also requested having an overview of such debugging information. This need is consistent with the findings of previous research [4, 16].

## 4.2 System Overview

The list of requirements have motivated us to develop a visual analytics system, CNNVis. It consists of the following components:

- A DAG formulation module to convert a CNN to a DAG and to aggregate neurons and layers for an overview (**R1,R4**);
- A neuron cluster visualization module to disclose the multiple facets of each neuron (**R3**);
- A biclustering-based edge bundling to reduce visual clutter caused by a large number of connections (**R4**);
- An interaction module that provides a set of interactions such as interactive clustering result modification (**R2**) and showing debug information on demand (**R5**).

The primary goal of CNNVis is to help experts better understand, diagnose, and refine CNNs. Fig. 4 illustrates the major components needed to achieve this goal. CNNVis takes a trained CNN and the corresponding training data set as the input. The input CNN is formulated as a DAG with each node representing a neuron and each edge representing the connection between neurons. To effectively present a large CNN, the DAG formulation module clusters the neurons in each layer. The clustered DAG is then passed to the neuron cluster visualization module. This module employs a rectangle packing algorithm to show the learned features of each neuron in a cluster and a matrix visualization to depict the activations of neurons. After that, a biclustering-based edge bundling clusters the edges to reduce visual clutter. Users can also interact with the generated visualization for further analysis. For example, users can interactively modify the neuron clustering results or show the average gradient of a selected layer.




Fig. 5. Illustration of the DAG formulation.

## 5 DAG FORMULATION

A CNN can be formulated as a DAG, where each node represents a neuron and each edge represents the connection between neurons. To effectively present a large CNN with tens or hundreds of layers and thousands of neurons in each layer, we first aggregate adjacent layers into groups. There are several ways to do the aggregation. For example, we can classify layers by merging two adjacent convolutional layers that have a small difference between their activation variance. We can also divide layers into groups at each pooling layer. In our current implementation, we employ the second one. In addition, the experts are interested in the output of an activation layer instead of that of a convolutional layer. As the outputs of these two layers have a one-to-one mapping relationship, we then merge these two layers and simply show the output of the activation layer (Fig. 5).

Then we cluster the neurons in each layer, which aims to group neurons with similar roles together. We assume that neurons with similar activations have similar roles. Directly using these activations to cluster the neurons is very time-consuming as there can be millions of images in the training set. Thus, we aggregate the activations into an average activation vector over the set of classes in the training set.

In particular, suppose the training samples can be categorized into  $m$  classes:  $c_1, c_2, \dots, c_m$ . The training samples of class  $c_i$  is represented by:  $S_i = \{s_1^{(i)}, s_2^{(i)}, \dots, s_{N_i}^{(i)}\}$ , where  $N_i$  is the number of training samples in class  $c_i$ . We first process each training sample  $s_j^{(i)}$  through the network and obtain the activation of neuron  $n$ :  $a_n(s_j^{(i)})$ . Then we calculate the average activation  $\bar{a}_n(c_i)$  of neuron  $n$  on class  $c_i$  by:

$$\frac{1}{N_i} \sum_{j=1}^{N_i} a_n(s_j^{(i)}). \quad (5)$$

Next, we combine each average activation into an activation vector  $\bar{a}_n = [a_n(c_1), a_n(c_2), \dots, a_n(c_m)]$ , which is a  $m$  dimension real-valued vector.

Finally, we cluster the neurons based on the derived activation vectors. In CNNVis, we employ two widely used clustering methods, K-Means [36] (parametric clustering) and MeanShift [8] (nonparametric clustering). The second method does not require prior knowledge of the cluster number. Thus, it is applicable to the case where experts do not know the cluster number of neurons. To better present each neuron cluster, we select several representative neurons that are closer to the cluster centroid.

## 6 VISUALIZATION

### 6.1 Overview

Based on the DAG formulation, we have designed a hybrid visualization (Fig. 6) that visually illustrates neuron clusters (nodes) and the connections between neurons (edges).

Each neuron cluster is represented by a large rectangle (Fig. 6A), which can be analyzed from multiple facets, such as the learned features,



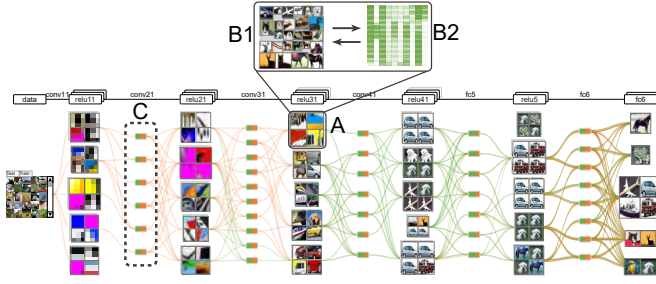


Fig. 6. Visualization overview.

activations, and contributions to the final result (**R3**). Specifically, we have adopted a rectangle packing algorithm to place the learned features of neurons in a neuron cluster, where each learned feature is encoded by a smaller rectangle (Fig. 6B1). Neuron activations are visualized as a matrix visualization (Fig. 6B2). Users can switch between the rectangle packing representation and the matrix visualization to explore different facets of the neurons.

To reduce visual clutter caused by dense edges and their crossings, we have developed a biclustering-based edge bundling algorithm (**R4**). For each layer, we first generate the biclusters between the input neuron clusters and output neuron clusters. Inspired by BiSet [48], we have also added an “in-between” layer between the input neuron clusters and output neuron clusters (Fig. 6C). In this layer, each bicluster is treated as a node in the DAG and is represented by a small rectangle.

In CNNVis, we employ the layout algorithm in TextFlow [9] to calculate the position of each node (e.g., neuron cluster or a bicluster) (**R1**). We also provide a set of interactions to facilitate deep analysis of a deep CNN (**R2**, **R5**).

Next, we will introduce the neuron cluster visualization and biclustering-based edge bundling in details.

## 6.2 Neuron Cluster Visualization

### 6.2.1 Learned Features as Rectangle Packing

**Computing learned features of neurons.** We employ the method used in [15] to compute the learned feature of a neuron because it is fast and the results are easier to understand. We also compute the activations of each neuron on a large set of image patches (e.g., sampled from the training set) and sort the patches in decreasing order according to their activations. To help experts better understand the role of each neuron, we select the top-5 patches with the highest activation scores to represent the learned feature of that neuron. By default, we show the top patch for a neuron and allow users to switch among these five patches. Other methods for computing the learned feature [35, 58] can easily also be integrated into CNNVis.




Fig. 7. Illustration of hierarchical rectangle packing.

**Layout.** A straightforward way to visualize the learned features (image patches) is to employ a grid-based layout where each image patch is represented by a rectangle of the same size [57, 58]. However, this method fails to emphasize the important neurons.

To tackle this issue, we formulate the layout of image patches as a rectangle packing problem, aiming to pack the given rectangles into an enclosing rectangle of a minimum area. We use the size of an image patch to encode the importance of the corresponding neuron because size is among the most effective visual channels [38]. In CNNVis, we provide several options to define the importance of a neuron, including its average or maximal activation on a set of classes and its contribution to the final result [33].

Existing rectangle packing algorithms [21, 28] can handle a small number of rectangles well (e.g., 15 rectangles in less than 0.1s [21]). However, the computing time grows exponentially as the number of packed rectangles increases (e.g., 25 rectangles in more than one hour [21]). Since a neuron cluster may consist of hundreds or even thousands of neurons, existing rectangle packing algorithms cannot directly be applied to our visualization.

To solve this problem, we have developed a hierarchical rectangle packing algorithm. The basic idea of our algorithm is to divide the problem into a number of smaller sub-problems. Each sub-problem can be efficiently solved by the state-of-the-art rectangle packing algorithm [21]. Specifically, our algorithm contains the following steps (Fig. 7).

**Step 1: Hierarchical clustering.** In this step, we perform a hierarchical clustering to divide the problem into several sub-problems that can be efficiently solved by an algorithm developed by Huang and Korf [21]. Specifically, we start with the cluster containing all of the neurons. Then we repeatedly split a cluster until the number of neurons in it is smaller than a threshold. This cluster splitting is done with a widely used graph clustering method [40].

**Step 2: Computing the layout area for each cluster.** Based on the hierarchical clustering result, we compute the layout area for each sub-cluster using a Treemap layout algorithm [26].

**Step 3: Rectangle packing of each cluster.** In this step, we compute the position and size for each image patch using the state-of-the-art rectangle packing algorithm [21].

### 6.2.2 Activations as Matrix Visualization

In our first prototype, we simply encode the activation of a neuron according to its size. However, the experts were not satisfied with that design because it failed to help them compare the roles of the neurons for different classes of images. To allow experts to compare different neurons, we stack the average activation vectors of neurons into an activation matrix, where each row is an average activation vector of a neuron. Accordingly, a matrix visualization is employed to visually illustrate the activation of the neurons. In particular, the color of a cell in the  $i$ -th row and  $j$ -th column represents the average activation of the  $i$ -th neuron  $n_i$  in class  $c_j$ .

This design was then presented to experts for evaluation. Overall, they liked the matrix visualization that provides a global overview of the activations among different classes. Their major concern was that the current visualization cannot reveal the cluster patterns in the activations of a neuron cluster. To solve this problem, we developed a matrix reordering algorithm that can visually reveal cluster patterns within the data.




Fig. 8. Illustration of matrix reordering: (a) before reordering; (b) after reordering.

**Matrix Reordering.** The order of columns (classes) should be consistent in different neuron clusters. Otherwise, experts are unable to directly compare the roles of neurons in two neuron clusters because of the different order of classes (columns). As a result, we only reorder the rows (neurons) in the matrix.

The basic idea of our algorithm is to maximize the sum of the similarities between adjacent neurons in the matrix. It aims to place neurons with similar activations close to each other, and thus can reveal the cluster pattern in the neuron cluster. Given neuron cluster  $C = \{n_1, n_2, \dots, n_{N_C}\}$ , the goal of the reordering is to find a row index  $\pi(i)$  for each neuron  $n_i$ , to better reveal the cluster pattern in a neuron cluster. For row  $r$  in the matrix, we denote its corresponding neuron

as  $n_{\pi^{-1}(r)}$ . To achieve this goal, we try to maximize the sum of the similarities between adjacent neurons in the matrix:

$$\max \sum_{r=1}^{N_C-1} \text{sim}(n_{\pi^{-1}(r)}, n_{\pi^{-1}(r+1)}), \quad (6)$$

where  $\text{sim}()$  is the similarity function between two neurons. In CNNVis, we adopt the widely used cosine similarity.

This combinational optimization problem can be solved by the Held-Karp algorithm [18] with a time complexity of  $O(2^{N_C} \cdot N_C^2)$ , where  $N_C$  is the number of neurons. The problem of directly applying it in our system is that we may have hundreds of neurons in a neuron cluster and the running time of the algorithm is very long. Thus, we developed a divide-and-conquer method to accelerate the algorithm, which consists of the following steps.

**Divide.** If the number of neurons in a cluster is too large to be efficiently solved via directly running the Held-Karp algorithm, the cluster is divided into several sub-clusters by a widely used graph clustering method developed by Newman [40].

**Conquer.** Computing the ordering of sub-clusters by running the Held-Karp algorithm.

**Combine.** Merging the ordering of sub-clusters into a global ordering.

Fig. 8 shows one result generated using our reordering method. With our method, several clusters can easily be detected.

### 6.2.3 Interaction

To better facilitate understanding of the multiple facets of each neuron cluster, CNNVis provides a set of user interactions.

**Interactive Clustering Result Modification.** Since the clustering algorithm is less than perfect and experts may have different needs, we allow experts to interactively modify the clustering results based on their knowledge (R2). Inspired by NodeTrix [19], we allow experts to drag a neuron out of a neuron cluster or to another neuron cluster.

**Selecting A Part of Neurons to View.** There are thousands of neurons in a CNN. Thus, it is necessary to allow experts to select some of the neurons to view. We allow users to select a set of classes and show the neurons that are strongly activated by the images in these classes. Other irrelevant neurons are deemphasized by setting them to be translucent.

**Switching between Facets.** Exploring the multiple facets of neurons can help experts better understand the roles of neurons. Thus, we allow users to switch between these facets (R3). For example, users can switch to view the learned features or the activation matrix.

### 6.3 Biclustering-based Edge Bundling

Initially, we visualized each edge as a curve. The major concern of the experts is visual clutter caused by millions of edges between nodes.

In order to reduce visual clutter, we tried two geometry-based edge bundling methods [11, 20] to cluster the edges between two layers. After interacting with CNNVis, the experts commented that this bundling method reduces visual clutter to some extent. However, the clusters revealed by the geometry-based bundling methods did not help their analysis because the edges with similar weights were not clustered together. The experts are more interested in edges with larger absolute weights, because this indicates that the corresponding inputs have a larger impact on the output.

To fulfill this requirement, we developed a biclustering-based edge bundling method to bundle edges with both similar and large absolute weights. For a given layer, a bicluster is a subset of input neuron clusters and a subset of output neuron clusters. This method can logically aggregate multiple individual connections and thus provides an opportunity to visually bundle edges between neuron clusters. Our algorithm contains the following steps (Fig. 9).

**Step 1: Aggregating Connections between Neurons.** We first calculate the strength  $w_{ij}$  of the connection  $e_{ij}$  between two neuron clusters,  $C_i$  and  $C_j$ . We denote  $E = \{e_{ij}\}$  as the edge set. An intuitive approach is to use the average of all the weights of the edges connecting a neuron  $n_s \in C_i$  and a neuron  $n_t \in C_j$ . The problem with this method is that it aggregates positive edges (edges with positive weights) and negative edges (edges with negative weights) and may result in an aggregated edge with a

small weight. This may lead to a misunderstanding. Thus, we calculate the strength of the connection between two neuron clusters as a two-dimensional vector  $\tilde{w}_{ij} = [w_{ij}^{pos}, w_{ij}^{neg}]$ , where  $w_{ij}^{pos}$  is the average of positive edge weights and  $w_{ij}^{neg}$  is the average of the negative edge weights.

**Step 2: Biclustering.** Based on the aggregation results, we then detect biclusters between the input neuron clusters and the output neuron clusters. Because experts are interested in both larger positive edges and smaller negative edges, we cannot simply convert it to an unweighted graph and perform biclustering. Thus, we first seek the maximum value  $w_{max}$  in  $W = \{w_{ij}^{pos}\} \cup \{|w_{ij}^{neg}|\}$ . If  $w_{max} \in \{w_{ij}^{pos}\}$ , then we select the edges satisfying:  $|w_{ij}^{pos} - w_{max}| < \tau$ , where  $\tau$  is a user defined parameter denoting the tolerance of similarity. If  $w_{max} \in \{|w_{ij}^{neg}|\}$ , we then perform the similar extraction. For these edges, we then mine the closed item sets as biclusters, where each input neuron cluster is connected to each output neuron cluster. To mine the closed item sets, we adopt the widely used Apriori, an algorithm for frequent item set mining [2]. After that, we remove the edges in the extracted biclusters from  $E$  and then repeat the process until  $w_{max}$  is under a user defined threshold.

**Step 3: Edge Bundling.** In this step, we bundled the edges in the same bicluster to reduce the visual clutter. Inspired by BiSet [48], we also add an “in-between” layer between the input neuron clusters and the output neuron clusters (Fig. 9 (c)). In this layer, each bicluster is visualized as a rectangle. In a bicluster, we use two colored regions (green and red) to indicate the proportion between the number of positive edges and of negative edges. An edge between two neuron clusters consists of two aggregated curves (Fig. 9A, and Fig. 9B), where green and red visually encode positive and negative weights, respectively. Since experts are less interested in analyzing edges with smaller absolute weights, they are not displayed by default. These edges can be shown per users’ request.

**Interaction.** The debugging information can help experts diagnose a failed training process. In CNNVis, we allow experts to analyze the debugging information at different granularities (R5). For example, they can change the color encoding of edges to analyze the gradient of each weight. Experts also have the option to view the average gradient at each layer as a line chart to get an overview of the debugging information.




Fig. 9. Illustration of biclustering-based edge bundling.

## 7 APPLICATION

In this section, we present the case studies to demonstrate how CNNVis help experts understand, diagnose and refine a CNN.

### 7.1 Overview

We have worked closely with the team of experts to select the base CNN model and to design the case studies.

**Base CNN.** The base CNN was contributed by  $E_3$  of the expert team. For brevity’s sake, we refer to the base CNN as BaseCNN. BaseCNN was designed based on a widely used deep CNN introduced in [46], which is often used in image classification. Recently, the expert team that we collaborate with has been redesigning this CNN and testing the performance of the variants. BaseCNN consists of 10 convolutional layers and two fully connected layers. The convolutional layers are organized into four groups, containing 2, 2, 3, and 3 convolutional layers, respectively. Each group is ended with a max-pooling layer. When designing BaseCNN, the expert employed the commonly used activation function, ReLU, and the commonly used loss function, cross-entropy. The architecture of BaseCNN is depicted in Fig. 10.

BaseCNN was trained and tested on a benchmark image dataset, CIFAR10 [29], which consists of 60,000 labeled color images of size  $32 \times 32$  in 10 different classes (e.g., airplane, bird, and truck), with 6,000 images per class. The dataset is split into a training set containing

50,000 images and a test set containing 10,000 images. Training and testing of BaseCNN are performed under a widely used deep learning framework, Caffe [25]. The BaseCNN model achieves 11.32% error on the test set.

**Design of Case Studies.** We have worked closely with the expert team to design three case studies from their current research on CNNs.

First, based on BaseCNN, the expert team constructed several variants and aimed to study the influence of the network architecture on the performance. The experts said that such an analysis would help to better understand the reason why CNNs with different architectures have different performance (Section 7.2).

Second, the expert team wanted to diagnose a training process that failed to converge. For example, in one training trial,  $E_3$  changed the output activation function and the loss function of BaseCNN. However, the training failed. The expert team wanted to diagnose the training process and find potential issues. This scenario triggered the second case study (Section 7.3).

Finally, the expert team wanted to further improve the performance of the BaseCNN model. To this end, the expert team decided to examine the output of each layer from a global overview to local details and detected a potential direction to improve the model. This requirement is addressed in the third case study.

Due to the page limit, we focus our report on the first two case studies. Interested readers may refer to the attached video for the study on model refinement (third case study).




Fig. 10. The architecture of BaseCNN. It contains four groups of convolutional layers and two fully connected layers. The number below a layer is the number of neurons in that layer.

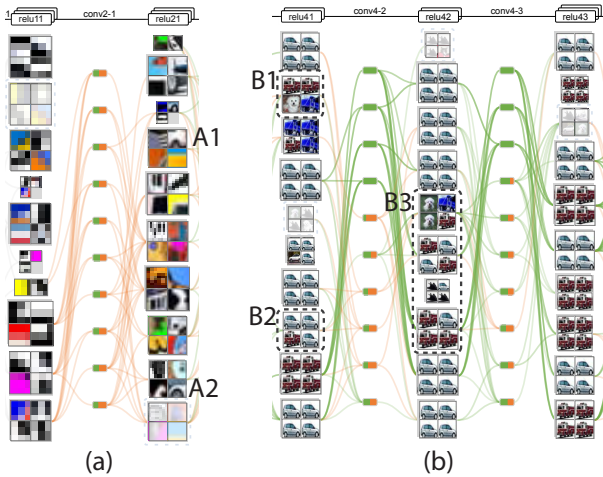


Fig. 11. Learned features of BaseCNN: (a) low level feature; (b) high level features.

## 7.2 Case Study: Influence of Network Architecture

This case study was a collaboration with expert  $E_2$ . In this case study,  $E_2$  evaluated the effectiveness of CNNVis on a set of variants of BaseCNN (with different depths and widths) qualitatively based on his experience. He also checked the possibility to select a CNN with a suitable architecture under the guide of CNNVis. Though a lot of high-performance models can be referred to on benchmark datasets, it usually takes a long

time to transfer the experience to other scenarios (e.g., choose a suitable CNN on a new dataset). Therefore,  $E_2$  emphasized that a systematic study on the network architecture and its influence on the performance is necessary to summarize reusable knowledge from existing trials and hopefully transfer it to the development process of other relevant deep models.

**Overview of BaseCNN.** We first provided expert  $E_2$  with an overview of BaseCNN (Fig. 1) to evaluate the quality of CNNVis.

From the overview, he identified that the neurons in the lower layers learned to detect simple patterns such as corners, color patches, and stripes (Fig. 1A). A similar observation was reported in previous work [30]. He identified a neuron for detecting a color patch in a cluster that mainly consists of neurons for detecting stripes with various orientations. To better compare the neurons that detect color patches, he dragged the neuron to a cluster that mainly consists of neurons for detecting color patches (Fig. 1B). Switching between the top-5 image patches that highly activate a given neuron in lower layers (Fig. 1A), he noticed that the retrieved patches did not show much difference in appearance. Then he turned to higher layers. After exploring among the top-5 image patches for a given neuron in higher layers (Fig. 1C), he noticed that these neurons could learn to detect more abstract features (e.g., an automobile). He concluded that, “The ability of detecting more abstract features in the higher layers is a nice property of well-trained deep CNNs and CNNVis indeed shows this pattern well.”

To further evaluate the ability of CNNVis to visualize the finer details of CNNs,  $E_2$  selected two similar classes (automobile and truck) and then examined the activation patterns of the relevant neurons. From the learned features in the lower layers, he found some common parts of trucks and automobiles, such as wheels (A1, A2 in Fig. 11 (a)). He indicated that these features are not sufficient to distinguish these two classes. Thus, he expanded the 4-th group of convolutional layers for further examination (Fig. 11 (b)). Expert  $E_2$  noticed that the number of “impure” neuron clusters gradually decreases as he moved to the higher layers. Here, an “impure” neuron cluster means that the image patches that maximally activate the neurons in the cluster are from different classes. Examining the “purity” means that we check the ability of a CNN to distinguish different semantics conveyed by class labels. In a pure cluster, the image patches that have the same semantics (class label) are gathered together in the activation space generated by the outputs of the layer. Note that in the lower layers, we prefer “impure” clusters because we want the neurons to detect as many different kinds of features as possible. While in higher layers, we prefer “pure” clusters because we want the model to separate higher-level semantics (different classes) by a large margin, so that the image patches from different classes seldom exist in the same cluster. We illustrate this criterion in Fig. 12. For example, in the top convolutional layer of BaseCNN, all clusters look “pure”, which indicates that the output activations given by BaseCNN match well with the semantics of different classes.




Fig. 12. Illustration of an “impure” cluster and a “pure” cluster.

**Network Depth.**  $E_2$  further investigated how the depth of the network affects the features detected by the neurons. He compared BaseCNN with two variant models, including ShallowCNN, which cuts off the 4-th group of convolutional layers, and DeepCNN, which doubles the number of convolutional layers. The architectures and accuracies are summarized in Table 1. He also selected the truck and automobile



classes, and expanded the last group of convolutional layers (Fig. 13 (a)). In ShallowCNN, he identified that there were indeed a lot more “impure” clusters in the top convolutional layers compared to those in BaseCNN, which indicates that a model without a sufficiently large depth is often incapable of distinguishing the images from similar classes, which can lead to a decrease of the performance. In DeepCNN, expert  $E_2$  noticed that almost all the weights in the first convolutional layer in the 4-th group were positive (Fig. 13 (b)). The expert commented that since the inputs of that layer were non-negative, the outputs are mostly positive. The outputs are then fed into ReLU. As ReLU retains a positive part of the inputs, the ReLU layer, together with its corresponding convolutional layer, can be viewed as a close-to-linear function. By further expanding the 4-th group of convolutional layers, expert  $E_2$  identified several consecutive layers that have a similar pattern (Fig. 14). Because the composition of linear functions is still linear, he concluded that this phenomenon indicates redundancy in the layers. He also commented that such redundancy may hurt overall performance and make the learning process computationally expensive and statistically ineffective. These findings are consistent with previous research [49].  $E_2$  then concluded that CNNVis could be used to check the abstractness of the features extracted by CNNs.

Table 1. Performance comparison between CNNs with different depth. “#ConvLayers” is the number of convolutional layers and “#Layers” is the number of layers that can be visualized.

	Error	#ConvLayers	#Layers
ShallowCNN	11.94%	7	30
BaseCNN	11.33%	10	40
DeepCNN	14.77%	20	70

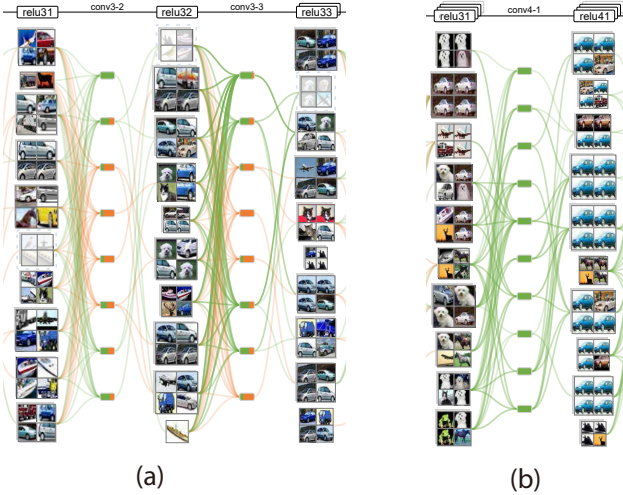


Fig. 13. Influence of the model depth: (a) high level features of a shallow CNN; (b) A convolutional layer whose weights are almost positive in DeepCNN.

**Network Width.** Another important factor that influences performance is the width of a CNN. To have a comprehensive understanding of its influence,  $E_2$  evaluated several variants of BaseCNN with different widths, named by BaseCNN $\times w$ , where  $w$  denotes the ratio of the number of neurons in a layer compared to that of BaseCNN. For example, BaseCNN $\times 4$  contains four times the neurons of BaseCNN. In the case study,  $w$  is selected from  $\{4, 2, 0.5, 0.25\}$ . The architecture and performance of these variants as well as BaseCNN are listed in Table 2.

Compared to BaseCNN, a wider network (BaseCNN $\times 4$ ) has a much lower training loss than testing loss. The expert commented that this phenomenon is known as overfitting in the field of machine learning. It means that the network tries to model every minor variation in the input, which is more likely to be noise. It often occurs when we have too many parameters relative to the number of training samples. When a model overfits, its performance on the testing set will be much worse than that

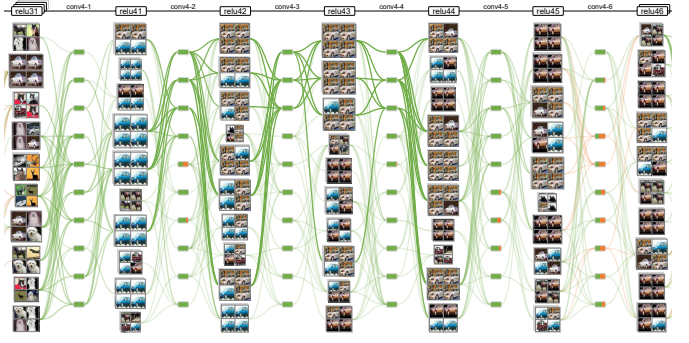


Fig. 14. Consecutive convolutional layers whose weights are almost positive in DeepCNN.

Table 2. Performance comparison between CNNs with different widths. #params is the number of parameters in the model, which is measured in millions.

	Error	#params	Training loss	Testing loss
BaseCNN $\times 4$	12.33%	4.22M	0.04	0.51
BaseCNN $\times 2$	11.47%	2.11M	0.07	0.43
BaseCNN	11.33%	1.05M	0.16	0.40
BaseCNN $\times 0.5$	12.61%	0.53M	0.34	0.40
BaseCNN $\times 0.25$	17.39%	0.26M	0.65	0.53

on the training set.  $E_2$  wanted to examine the influence of overfitting on CNNs. He visualized BaseCNN $\times 4$  with our visual analytics system.

After examining the higher level features, the expert did not found much difference compared to BaseCNN. Then he switched to examine low level features. He instantly found that several neurons learn to detect almost the same features (Fig. 15 (a)). The expert inferred that there may be redundant neurons in an overfitting CNN. For further verification, he decided to examine the activations of the neurons in this cluster. Compared to the activations in lower layers of BaseCNN (Fig. 15 (b)), he found that many neurons have very similar activations. This observation verified that there are redundant neurons in the lower layers of a CNN that is too wide.

$E_2$  commented, “We often use a quantitative criterion (e.g., accuracy) to evaluate the quality of a model. However, a quantitative criterion itself cannot provide sufficient intuition and clear guidelines. Even I know a CNN overfits, it is hard to decide which layer to narrow down or remove. While CNNVis can guide me to locate the candidate layers, which is very useful in my research.”

$E_2$  then compared the performance of BaseCNN with narrower networks (BaseCNN $\times 0.5$  and BaseCNN $\times 0.25$ ). Although the training loss and testing loss of these narrower networks are comparable, which indicate that these narrow networks generalize well, their performance was worse than BaseCNN (Table 2). The expert explained that this phenomenon is known as underfitting. It happens when the task is complex but we are trying to use a simple model to perform the task. In image classification, one of the major disadvantage of underfitting is that the model is too simple to distinguish images from similar classes (e.g., automobiles and trucks). In addition to the decrease in accuracy, he wanted to know the influence that underfitting brought to the model.

The expert visualized BaseCNN $\times 0.25$  for further exploration. He selected two similar classes, automobile and truck, to examine the patterns of the relevant neurons. After analyzing low level features, he did not find much difference compared to BaseCNN. Thus, he switched his attention to high level features. When examining the features of the last convolutional layer, he found that there were several “impure” neuron clusters. For example, cluster C in Fig. 15 (c) is represented by three trucks and an automobile (outlier). He switched to explore the activations in this cluster (Fig. 15 (c)). The expert found the outlier has similar activations on the two classes (i.e., truck and automobile), which means that this neuron can hardly distinguish automobiles from trucks. As a result, the ability of the model to correctly classify images from similar classes is hindered, which is reflected in the decrease of accuracy.

Expert  $E_2$  commented that, “It is really hard for me to choose the



architecture, including the depth and width of the network on a new dataset, as there are not many high-quality deep models to refer to. I usually need to try a series of parameters to achieve a satisfactory performance. CNNVis can intuitively show the quality of the model in various ways, such as the purity of clusters, and help me find the suitable architecture more quickly.”




Fig. 15. Comparison between models with different widths: (a) low level features of BaseCNN×4; (b) low level features of BaseCNN; (c) high level features of BaseCNN×0.25

### 7.3 Case Study: Training Diagnosis




Fig. 16. Exploring the connections between neurons: (a) edges encoded by the relative change of weights; (b) edges encoded by the weights.




Fig. 17. Exploring the neuron clusters.

This case study demonstrates how CNNVis helps an expert ( $E_3$ ) diagnose a failed training process. Recently, during the research triggered by [32],  $E_3$  tried to construct a variant of BaseCNN. Specifically, he replaced the output activation function with the identity function (i.e.,  $f(x) = x$ ) and the loss function with the hinge loss (see the loss function part in Sec. 3). However, the training of this model failed. The problem was that the training process got stuck when the loss decreased to around 2.0, where the model was far from achieving a good accuracy.

To help the expert diagnose the failed training process, we provided him with the visualization of a snapshot after the training process got stuck. As he often uses the relative changes of weights to diagnose a training process in his previous research, he set the initial color coding of edges as the relative changes of weights.

From the overview, expert  $E_3$  observed that the edges were difficult to recognize after the top-2 layers (Fig. 16(a)). This indicated that the relative changes of weights were very small, which caused the training process being stuck.  $E_3$  was curious about what led to such small relative changes in weights, so he used the color of edges to represent the weights. He immediately identified that an overwhelming majority of edges were negative (Fig. 16(b)).

He wanted to find what influence the negative weights had on the model. As the learned features could not reveal too much information due to the failed training process, expert  $E_3$  switched to examine the activation matrix. He spotted some neuron clusters where all the neurons had zero activations on all classes. To further study this phenomenon, he sequentially expanded the second, third, and fourth groups of convolutional layers. He found that the ratio of neurons with zero activations became larger and larger from the lower layers to the higher layers (Fig. 17). The activation functions of these neurons are ReLUs. He continued to zoom in and further examine the inputs fed into the ReLUs, which he found were always negative. If the input of a ReLU is less than zero, it generates a zero activation.

Expert  $E_3$  explained that because the input of each convolutional layer is the output of ReLUs in the previous layer, it must be nonnegative. As the weights of the linear transformation in this layer are mostly negative, the values fed into ReLUs are mostly negative. Consequently, the outputs of ReLUs are mostly zeros. In the training method that we used (i.e., stochastic gradient descent [6]), zero outputs of a neuron mean zero updates to its weights.

Having learned why the training process got stuck, expert  $E_3$  proposed a method to force the network away from that situation. He added a batch-normalization layer [22] after each convolutional and fully-connected layer, before the ReLU activation function. With batch-normalization, the input fed into the ReLUs should no longer be mostly negative. This means that the model could still be trained even most weights were negative.

The improved model achieved an average error of 9.43% on the CIFAR-10 dataset, with which expert  $E_3$  was very satisfied. He further commented, “I have investigated this problem for a long time and inserted all kinds of code fragments to print the debugging information during training. However, after many unsuccessful attempts and a great deal of effort spent reading the debugging information, I eventually gave up. It is awesome to have a toolkit like CNNVis, which intuitively illustrates the training statistics and allows me to explore the training process from multiple perspectives.”

## 8 CONCLUSION

In this paper, we have presented a novel visual analytics system to help machine learning experts better understand, diagnose, and refine CNNs. Powered by a hybrid visualization consisting of rectangle packing, matrix ordering, and biclustering-based edge bundling, the system allows experts to explore and understand a deep CNN from different perspectives. In addition, it enables experts to diagnose and refine the CNN architecture to further improve the performance. Three case studies were conducted to demonstrate the effectiveness and usefulness of the system for comprehensive analysis of CNNs.

There are several directions for future work to further improve our system. Currently, CNNVis focuses on analyzing a snapshot of the CNN model in the training process, which is useful for conducting the offline analysis. All the experts expressed the need to integrate CNNVis with the online training process and continuously get an update of the training status. A key issue is the difficulty of selecting representative snapshots and comparing them effectively.

Another interesting venue for future work is to apply CNNVis to other types of deep models that cannot be formulated as a DAG, such as recurrent neural network (RNN). The major bottleneck is to design an effective visualization to facilitate experts in understanding the data flow through different types of deep models. For example, in addition to the conventional multi-layer neural network, RNN has a feedback loop from an output to an input. Better understanding the working principle of the feedback loop help experts design more effective models.

## REFERENCES

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [2] R. Agrawal and R. Srikant. Fast algorithms for mining association rules in large databases. In *VLDB*, pages 487–499, 1994.

- [3] Y. Bengio. Learning deep architectures for ai. *Foundations and Trends in Machine Learning*, 2(1):1–127, 2009.
- [4] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE PAMI*, 35(8):1798–1828, Aug 2013.
- [5] J. Bergstra, O. Breuleux, F. Bastien, P. Lamblin, R. Pascanu, G. Desjardins, J. Turian, D. Warde-Farley, and Y. Bengio. Theano: a CPU and GPU math expression compiler. In *SciPy*, 2010.
- [6] L. Bottou. Stochastic gradient learning in neural networks. *Neuro-Nimes*, 91(8), 1991.
- [7] R. Collobert, S. Bengio, and J. Mariéthoz. Torch: a modular machine learning software library. Technical report, IDIAP, 2002.
- [8] D. Comaniciu and P. Meer. Mean shift: a robust approach toward feature space analysis. *IEEE PAMI*, 24(5):603–619, 2002.
- [9] W. Cui, S. Liu, L. Tan, C. Shi, Y. Song, Z. J. Gao, H. Qu, and X. Tong. Textflow: Towards better understanding of evolving topics in text. *IEEE TVCG*, 17(12):2412–2421, 2011.
- [10] W. Cui, S. Liu, Z. Wu, and H. Wei. How hierarchical topics evolve in large text corpora. *IEEE TVCG*, 20(12):2281–2290, 2014.
- [11] W. Cui, H. Zhou, H. Qu, P. C. Wong, and X. Li. Geometry-based edge clustering for graph visualization. *IEEE TVCG*, 14(6):1277–1284, 2008.
- [12] A. Dosovitskiy and T. Brox. Inverting visual representations with convolutional networks. *arXiv preprint arXiv:1506.02753*, 2015.
- [13] D. Erhan, Y. Bengio, A. Courville, and P. Vincent. Visualizing higher-layer features of a deep network. Technical report, University of Montreal, 2009.
- [14] S. Gad, W. Javed, S. Ghani, N. Elmqvist, T. Ewing, K. N. Hampton, and N. Ramakrishnan. Themedelta: Dynamic segmentations over temporal topic models. *IEEE TVCG*, 21(5):672–685, 2015.
- [15] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, pages 580–587, 2014.
- [16] X. Glorot and Y. Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, pages 249–256, 2010.
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [18] M. Held and R. M. Karp. A dynamic programming approach to sequencing problems. *SIAM*, 10(1):196–210, 1962.
- [19] N. Henry, J.-D. Fekete, and M. J. McGuffin. Nodetrix: a hybrid visualization of social networks. *IEEE TVCG*, 13(6):1302–1309, 2007.
- [20] D. Holten and J. J. Van Wijk. Force-directed edge bundling for graph visualization. In *CGF*, volume 28, pages 983–990, 2009.
- [21] E. Huang and R. E. Korf. Optimal rectangle packing: An absolute placement approach. *Journal of Artificial Intelligence Research*, 46:47–87, 2012.
- [22] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456, 2015.
- [23] S. Jänicke, J. Focht, and G. Scheuermann. Interactive visual profiling of musicians. *IEEE TVCG*, 22(1):200–209, 2016.
- [24] K. Jarrett, K. Kavukcuoglu, M. Ranzato, and Y. LeCun. What is the best multi-stage architecture for object recognition? In *ICCV*, pages 2146–2153, 2009.
- [25] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [26] B. Johnson and B. Shneiderman. Tree-maps: A space-filling approach to the visualization of hierarchical information structures. In *Visualization*, pages 284–291, 1991.
- [27] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, pages 1725–1732, 2014.
- [28] R. E. Korf, M. D. Moffitt, and M. E. Pollack. Optimal rectangle packing. *Annals of Operations Research*, 179(1):261–295, 2010.
- [29] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Montreal, 2009.
- [30] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [31] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [32] C. Li, J. Zhu, T. Shi, and B. Zhang. Max-margin deep generative models. In *Advances in Neural Information Processing Systems*, pages 1828–1836, 2015.
- [33] H. Li, T. Jiang, and K. Zhang. Efficient and robust feature extraction by maximum margin criterion. *Neural Networks*, 17(1):157–165, 2006.
- [34] S. Liu, Y. Wu, E. Wei, M. Liu, and Y. Liu. Storyflow: Tracking the evolution of stories. *IEEE TVCG*, 19(12):2436–2445, 2013.
- [35] A. Mahendran and A. Vedaldi. Understanding deep image representations by inverting them. In *CVPR*, pages 5188–5196, 2015.
- [36] S. Marsland. *Machine learning: an algorithmic perspective*. CRC press, 2015.
- [37] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [38] T. Munzner. *Visualization Analysis and Design*. CRC Press, 2014.
- [39] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, pages 807–814, 2010.
- [40] M. E. Newman. Fast algorithm for detecting community structure in networks. *Physical review E*, 69(6):066133, 2004.
- [41] A. Nguyen, J. Yosinski, and J. Clune. Multifaceted feature visualization: Uncovering the different types of features learned by each neuron in deep neural networks. *arXiv preprint arXiv:1602.03616*, 2016.
- [42] A. r. Mohamed, G. E. Dahl, and G. Hinton. Acoustic modeling using deep belief networks. *IEEE TASLP*, 20(1):14–22, 2012.
- [43] F. Seide, G. Li, and D. Yu. Conversational speech transcription using context-dependent deep neural networks. In *Interspeech*, pages 437–440, 2011.
- [44] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.
- [45] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR Workshop*, 2013.
- [46] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [47] G. Sun, Y. Wu, S. Liu, T. Q. Peng, J. J. H. Zhu, and R. Liang. Evoriver: Visual analysis of topic competition on social media. *IEEE TVCG*, 20(12):1753–1762, 2014.
- [48] M. Sun, P. Mi, C. North, and N. Ramakrishnan. Biset: Semantic edge bundling with biclusters for sensemaking. *IEEE TVCG*, 22(1):310–319, 2016.
- [49] S. Sun, W. Chen, L. Wang, X. Liu, and T.-Y. Liu. On the depth of deep neural networks: A theoretical view. In *AAAI*, 2016.
- [50] Y. Tanahashi, C. H. Hsueh, and K. L. Ma. An efficient framework for generating storyline visualizations from streaming data. *IEEE TVCG*, 21(6):730–742, 2015.
- [51] Y. Tanahashi and K. L. Ma. Design considerations for optimizing storyline visualizations. *IEEE TVCG*, 18(12):2679–2688, 2012.
- [52] F. Y. Tzeng and K. L. Ma. Opening the black box - data driven visualization of neural networks. In *IEEE VIS*, pages 383–390, 2005.
- [53] D. Wei, B. Zhou, A. Torralba, and W. Freeman. Understanding intra-class knowledge inside cnn. *arXiv preprint arXiv:1507.02379*, 2015.
- [54] K. Wongsuphasawat and D. Gotz. Exploring flow, factors, and outcomes of temporal event sequences with the outflow visualization. *IEEE TVCG*, 18(12):2659–2668, 2012.
- [55] Y. Wu, N. Pitipornvivat, J. Zhao, S. Yang, G. Huang, and H. Qu. egoslider: Visual analysis of egocentric network evolution. *IEEE TVCG*, 22(1):260–269, 2016.
- [56] P. Xu, Y. Wu, E. Wei, T. Q. Peng, S. Liu, J. J. H. Zhu, and H. Qu. Visual analysis of topic competition on social media. *IEEE TVCG*, 19(12):2012–2021, 2013.
- [57] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, and H. Lipson. Understanding neural networks through deep visualization. In *ICML Workshop on Deep Learning*, 2015.
- [58] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*, pages 818–833, 2014.
- [59] M. D. Zeiler, G. W. Taylor, and R. Fergus. Adaptive deconvolutional networks for mid and high level feature learning. In *ICCV*, pages 2018–2025, 2011.