

VAC-CNN: A Visual Analytics System for Comparative Studies of Deep Convolutional Neural Networks

Xiwei Xuan, Xiaoyu Zhang, Oh-Hyun Kwon, and Kwan-Liu Ma

Abstract—The rapid development of Convolutional Neural Networks (CNNs) in recent years has triggered significant breakthroughs in many machine learning (ML) applications. The ability to understand and compare various CNN models available is thus essential. The conventional approach with visualizing each model’s quantitative features, such as classification accuracy and computational complexity, is not sufficient for a deeper understanding and comparison of the behaviors of different models. Moreover, most of the existing tools for assessing CNN behaviors only support comparison between two models and lack the flexibility of customizing the analysis tasks according to user needs. This paper presents a visual analytics system, VAC-CNN (Visual Analytics for Comparing CNNs), that supports the in-depth inspection of a single CNN model as well as comparative studies of two or more models. The ability to compare a larger number of (e.g., tens of) models especially distinguishes our system from previous ones. With a carefully designed model visualization and explaining support, VAC-CNN facilitates a highly interactive workflow that promptly presents both quantitative and qualitative information at each analysis stage. We demonstrate VAC-CNN’s effectiveness for assisting novice ML practitioners in evaluating and comparing multiple CNN models through two use cases and one preliminary evaluation study using the image classification tasks on the ImageNet dataset.

Index Terms—Visual analytics, machine learning, convolutional neural network, model comparison, visual explanation.



1 INTRODUCTION

IN recent years, researchers have pushed the boundary of various domains unprecedentedly by taking advantage of the state-of-the-art deep Convolutional Neural Networks (CNNs) [1], [2], [3], [4], [5], [6], [7], [8], [9]. During this process, many machine learning (ML) practitioners with diverse knowledge backgrounds share the common need to understand and compare multiple CNNs. Such comparison tasks are challenging for novice ML practitioners who have primary but not comprehensive ML knowledge background, especially when the number of models to compare is large and the features of them vary a lot. For example, a medical school graduate student may want to adopt a CNN for disease detection. With tens of different CNN architectures available, it is difficult for them to filter out inapplicable models, let alone to find one with desired features. Conventional approaches for comparing multiple CNNs [10], [11], [12], [13], [14] often focus on investigating model architectures [10], [11], [12] or analyzing quantitative performances statically [10], [13], [14], but fail to provide enough intuitive information or reasons behind the different behavior of models. Therefore, it calls for efforts to develop novice-friendly tools for ML practitioners that improve models’ transparency, reveal models’ differences, and extend models’ applications through understanding their behaviors in CNN comparative studies.

An interpretable CNN comparative study can be divided into two phases—model interpretation and model comparison. For model interpretation, researchers from the XAI (eXplainable Artificial Intelligence) [15] community have developed plenty of class-discriminative visual explanation methods as a post-hoc

analysis of the underlying behaviors of deep models [16], [17], [18], [19], [20], [21]. These methods highlight the region of interest (ROI) relevant to the model’s decision, and could significantly increase the interpretability of deep models [16]. However, most of them are only applied to analyze a single model’s behaviors in detail, while rarely used to compare multiple models. For model comparison, many visual analytics tools have been developed for interactive CNN comparison [22], [23], [24], [25], [26], [27]. They integrate different visualization techniques to compare deep models from different perspectives, such as feature activations, parameter distributions, etc. Some of these tools support multi-model comparison, but they either lack interpretability [22], [23], [24], [25] or only support comparison between two models [23], [26], [27]. In response to the increasing number of models to compare and choose from, it is necessary to consolidate the state-of-the-art techniques from both phases and develop a CNN model comparative study tool that can take a flexible number of models and provide explanations for model behavior.

In this paper, we introduce a visual analytics system—VAC-CNN (Visual Analytics for Comparing CNNs)—to support an interpretable comparative study of deep CNNs. VAC-CNN assists the progress of a highly interactive workflow with carefully designed visualizations. To facilitate flexible comparison customization, VAC-CNN supports three types of comparison studies: 1) high-level screening for a large number of (e.g. tens of) models, 2) behavior consistencies evaluation for a few models, and 3) detailed investigation for a single model. To enhance models’ interpretability, VAC-CNN integrates multiple class-discriminative visual explanation methods, including Grad-CAM [17], BBMP [18], Grad-CAM++ [19], Smooth Grad-CAM++ [20], and Score-CAM [21]. To present the results of these methods smoothly, VAC-CNN promptly visualizes both quantitative and qualitative information

- X. Xuan, X. Zhang, O.-H. Kwon and K.-L. Ma are with University of California, Davis.
E-mail: {xwxuan, xybzhang, kw, klma}@ucdavis.edu.

Manuscript received October 22, 2021; revised January 14, 2022.

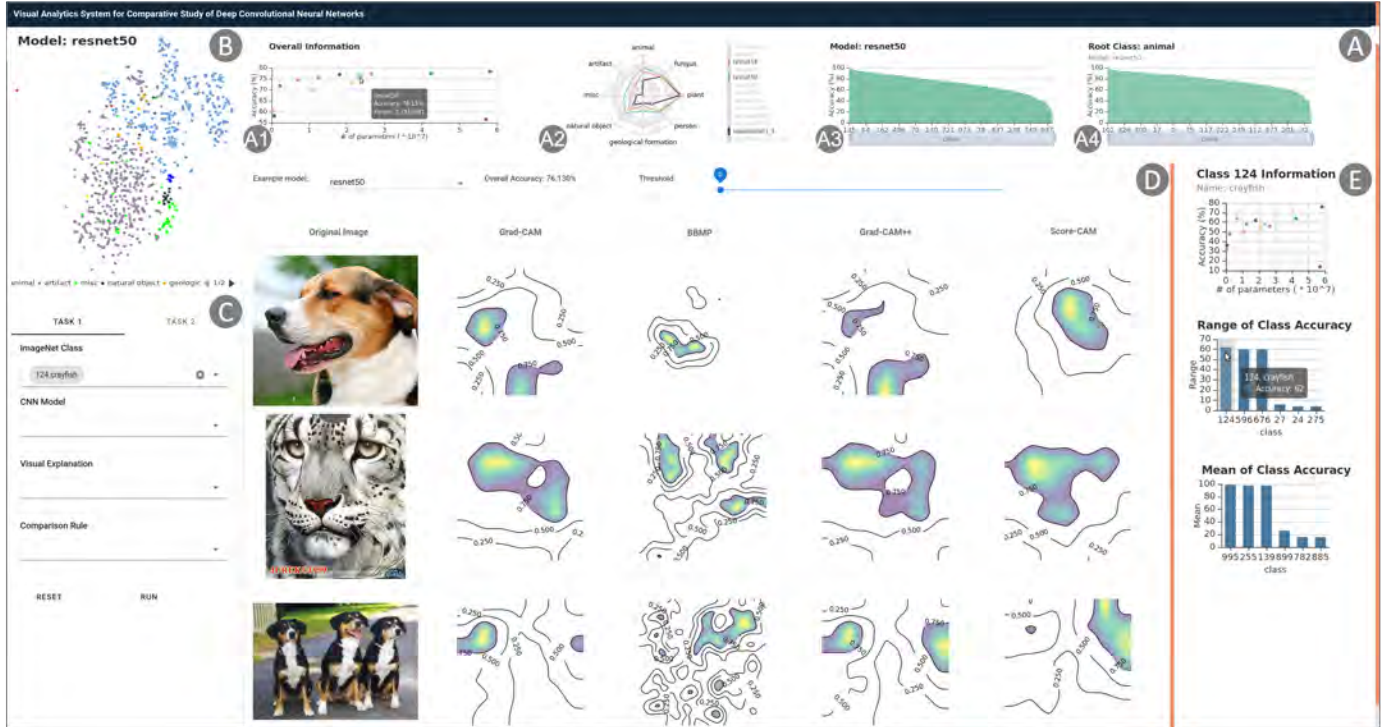


Fig. 1. The interface of our visual analytics system. (A) **Overall Information View** visualizes multiple quantitative matrices of each CNN model. (B) **Distribution Graph View** shows the aggregating activations of 1000 ImageNet classes generated from the distance matrix of each model’s predictions. (C) **Task Selection Sidebar** facilitates users to customize multiple selections to interactively compare the models’ behaviors. (D) **Visual Explanation View** provides two types of information - before running any task, it presents the visual explanation examples; when users are running a task, it shows the explanation results. (E) **Supplemental View** works as a supplementary information board.

at each analysis stage, allowing users to investigate and compare multiple models from different perspectives.

We illustrate the effectiveness of our visualization and interaction design to assist ML novices in CNN interpretation and comparison with two use cases. One is about multi-model comparison on a single input image, and the other is about single model behavior inspection on different classes of images. We also evaluate the usefulness of VAC-CNN with a preliminary evaluation study. According to the evaluation result, our system is easy to use and capable of providing useful insights about model behavior patterns for novice ML practitioners.

The primary contributions of our work include:

- A visual analytics system to support flexible CNN model analysis from single-model inspection to multi-model comparative study.
- A suite of enhanced visual explanation methods coordinated by a highly interactive workflow for effective and interpretable model comparison.

2 RELATED WORK

Our system for the comparative study of interpretable CNN models is inspired by previous works related to deep learning and XAI. This section discusses existing research on visual explanation methods for understanding CNN model behaviors, CNN model comparison, and visual analytics for interpretable CNN comparison.

2.1 Visual Explanation Methods for Interpretable CNN

Visual explanation methods play an essential role in improving the transparency of deep CNN models. According to the visualization purpose, the existing visual explanation methods can be grouped

into three kinds. The first group of methods mainly focus on visualizing the activations of neurons and layers inside a specific model, such as Feature Visualization [28] and Deep Dream [29]. These methods focus on exploring a single model’s internal operation mechanism, which is not scalable for comparing multiple models. The second group of methods represents the view of an entire model, which visualizes all extracted features of a model without highlighting decision-related information, such as Vanilla Backpropagation [30], Guided Backpropagation [31], and Deconvolution [32]. This group of explanations’ primary processing method is the backward pass, which is time-efficient and can produce fine-grained results. However, this group of methods fails to explain models’ decision-making convincingly because they indistinctively represent all the extracted information.

The third group of methods is the class-discriminative visual explanation [16], [17], [18], [19], [20], [21], which can explain the model decision by localizing the regions essential for model predictions and is sensitive to different classes. Zhou et al. [16] introduce CAM (Class Activation Map), which is an initial approach of localizing a specific image region for a given image class. However, researchers have to re-train the entire model to get the results of CAM. As an approach to address the shortcomings of CAM, Grad-CAM [17] is proposed as a more efficient approach, which can explain the predictions of CNN models without re-training or changing their structure. In 2017, a perturbation-based method called BBMP [18] was introduced, which highlights the ROI of input images with the help of perturbations on input images. Since BBMP requires additional pre-processing and multiple iterations, it was time-consuming and challenging to be implemented in real-time applications. Recently, plenty of Grad-CAM-inspired methods

have been proposed including Grad-CAM++ [19], Smooth Grad-CAM++ [20], and Score-CAM [21]. Consistent with Grad-CAM, these methods are applicable to a wide variety of CNN models.

Aiming to provide interpretable CNN model comparison, we include multiple class-discriminative visual explanation methods to support the understanding of models' decisions.

2.2 CNN Model Comparison

Lots of previous works are aiming to address the need for CNN model comparison. To assist researchers in CNN model evaluation and comparison, Canziani et al. [13] develop a quantitative analysis of fourteen different CNN models based on the accuracy, memory footprint, number of parameters, operations count, inference time, and power consumption. In terms of statistical analysis, multiple findings concerning the relationships of model parameters are discussed in [13], such as the independence between power consumption and architecture, the hyperbolic relationship of accuracy, and inference time. Liu et al. [11] go through four kinds of deep learning architectures, including autoencoder, CNN, deep belief network, and restricted Boltzmann machine. It also illustrates those architectures' applications in some selected areas such as speech recognition, pattern recognition, and computer vision. A recent survey by Khan et al. [12] discusses the architecture development of deep CNNs, from LeNet [2] presented in 1998 to Comprehensive SqueezeNet [9] presented in 2018. [12] offers a detailed quantitative analysis of twenty-four deep CNN models, comparing information such as the number of parameters, error rate, and model depth.

Besides these general model comparison studies [11], [12], [13], researchers also apply the comparative study of multiple models for specific tasks. Aydogdu et al. [10] quantitatively compare three different CNN architectures based on their performance in the age classification task. Talebi et al. [33] train multiple models to automatically assess image quality and compare their performances based on the accuracy and other quantitative measurements. Mukhopadhyay et al. [14] apply the performance comparison of three CNN models for the Indian Road Dataset, which represents the road detection results through images and compares the models based on the detection accuracy. By discussing past model comparison studies, we conclude that conventional works focus on quantitative comparison or structure analysis, which fail to reveal the underlying reasons for models' performances. To fill this research gap, our system would integrate XAI techniques, specifically, the visual explanation methods, to help researchers compare the deep CNN models in an interpretable way.

2.3 Visual Analytics for Interpretable CNN Comparison

A variety of visual analytics tools are aiming at supporting interpretable comparisons of CNN. Some focus on visualizing and interpreting the internal working mechanism of a single CNN model [34], [35], [36], [37], [38], [39], [40], combining various visualization techniques such as dimension reduction for understanding networks' hidden activities [34], a directed acyclic graph to disclose multiple neurons' facets and interactions [35], hierarchy analysis of similar classes [37], or feature visualizations and interactions [39]. However, such in-depth inspection of a single model helps develop interpretation but is insufficient for scenarios where model comparison and selection are needed.

Researchers have developed some visual analytics frameworks for comparing CNN models [22], [23], [24], [25], [26], [27], [41], [42]. Prospector [22] leverages partial dependence plots

to visualize different performances of multiple models on one input sample. To assist model training, CNN Comparator [23] compares models from different training stages in the aspect of model structures, parameter distributions, etc. Utilizing predictions of the labels, Manifold [24] allows users to compare multiple models at the feature level using scatter plots. BEAMES [25] is a multimodel steering system providing multi-dimensional inspection to help domain experts in model selection. However, these methods lack interpretability because they mainly use numerical features of CNNs. To assist interpretable comparison, researchers apply techniques such as linking model structures instances for comparing two binary classifiers [26], visualizing qualitative differences in how models interpret input data [27], etc. These techniques can assist better model interpretation, but only support the comparison among a small number of models.

In conclusion, most of the existing visual analytics methods for interpretable CNN comparison are either based on handcrafted quantitative parameters or only supporting comparison between two models. Only few of them allow CNN interpretation and multi-model comparison at the same time. With comparing and interpreting different CNN models becoming a growing demand, there is a need for comparative studies that support a larger number of models' comparisons and present quantitative and qualitative information at the same time for more thorough evaluations.

3 DESIGN GOALS

According to our survey, we are aware of the need for CNN comparison tools that support flexible customization of comparative tasks (e.g., the in-depth inspection of a single model and comparative studies of multiple models). Such tools should also integrate XAI techniques to assist model interpretation. We refine this requirement into four design goals and describe them as follows.

G1 Novice-Friendly Information Overview: Motivated by the superb learning power of CNN models, researchers from different domains with various knowledge backgrounds are attempting to take advantage of this fast-developing technique in recent years [43]. A visual analytics system for CNN comparison can be helpful for beginners as well as experts to gain more insights on models' behaviors. Given that most of the existing model comparison tools are developed for experienced ML researchers, our system needs to provide an information overview that can assist users in a high-level model screening based on their performances and a general understanding of the XAI techniques we integrated. Moreover, the system should distill information and enable interactions to assist the overview process instead of overwhelming users with too many details all at once.

G2 Informative Visual Explanation: The commonly employed visual explanation methods based on color heatmaps highlighting the associated ROI are shown to be helpful in interpreting CNNs [16], [17], [18], [19], [20], [21]. However, it is hard to efficiently identify differences among models only based on the qualitative results from such visual explanation methods in the model comparison scenarios. Thus, we need to consolidate the visual explanation methods with quantification measurements to help users gain better insights during the comparison process. Besides, when localization is not enough for interpreting a model, our system should provide complementary visualization for further analysis and help

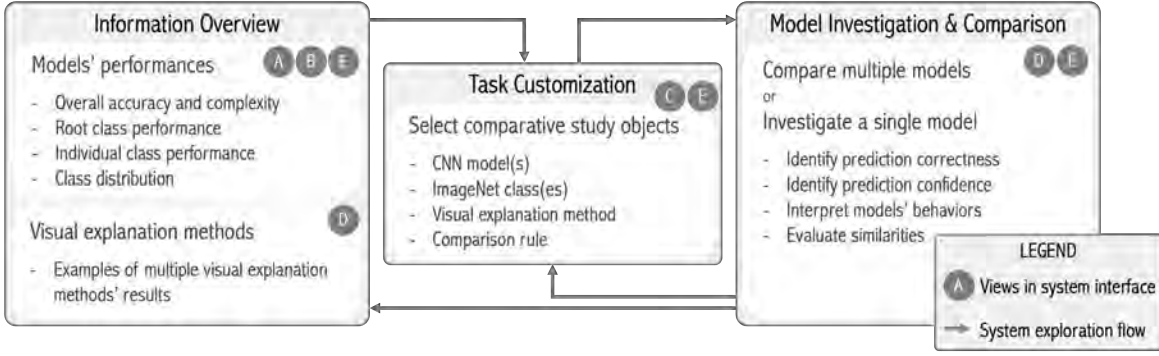


Fig. 2. The comparative analysis workflow of VAC-CNN consists of three phases: information overview, task customization, and model investigation & comparison. Each phase involves user interaction with multiple views of our system.

users better understand the underlying reason behind the CNN model predictions.

G3 Scalability and Flexibility: Unlike ML experts, beginners without comprehensive ML knowledge can benefit from additional exploration in a broader range of models when comparing models. Therefore, they need a visual analytics tool that supports scalability in the number of models to compare and flexibility in the customization of comparison tasks [24]. However, most of the existing comparison approaches for analyzing model behaviors only focus on two-model comparison [23], [26]. To fill in this gap, we need to support scalable and flexible CNN comparison tasks in our system and allow users to customize objects such as the number of the model(s), data class(es), and the visual explanation method(s).

G4 Real-time Interaction: It could take a tremendous amount of GPU time to generate models' visual explanation results [18], [21], especially for large-scale datasets. With a web-based approach, we expect our system to be efficient enough to offer users a responsive interface, which means users should not experience a noticeable delay when exploring model comparison scenarios through our system. Besides, we should allow users to interactively audit details of each view and select specific elements to inspect further information. Moreover, it is essential to present multiple views synergistically and help users better understand the models through the coordinated information of each view.

4 METHODOLOGY

VAC-CNN is built upon thirteen widely-used models, to cover various state-of-the-art architectures such as AlexNet [3], ResNet [4], SqueezeNet [5], DenseNet [6], MobileNet [7], and ShuffleNet [8]. The models are pre-trained on the ImageNet dataset [44] for the image classification task and we develop our system on the ImageNet (ILSVRC2012) validation set with 1,000 image classes and 50,000 images. In this section, we introduce the analysis workflow and the integrated methodologies of our system.

4.1 Workflow

Based on our survey of existing tools [22], [23], [24], [25], [26], [27], [41], [42] and the design goals in Sec. 3, we model the comparative analysis procedure with VAC-CNN as a three-phase workflow (see Fig.2). The workflow starts from **Phase 1** which provides an information overview to help ML beginners get a general

understanding of both model performances and visual explanation methods. **Phase 2** provides task customization to support flexible study options towards CNNs, ImageNet classes, visual explanation method, and comparison rule. Based on the customized comparison requirements, **Phase 3** presents coordinated visualizations and qualitative information for multi-model comparison or single-model investigation, respectively. We will connect our discussion about the methodology in this section and the interface design in the following section with these phases.

4.2 Distribution Graph Generation

In regard to design goal **G1**, we provide a comprehensive and novice-friendly information overview for our users to understand the model's high-level performances. One way to achieve this is to investigate the class distribution which is generated from the model's prediction and reflects how the model interprets the data.

To visualize the distribution of image classes in respect to a specific model, we create a distribution graph (Fig. 1 (B) based on each model's predictions. In Algorithm 1, we use confidence matrix as the baseline to generate this graph, since it reflects how a model understands the input data. For a given model, every input image is classified according to a confidence vector of size $(1, N)$, where N is the number of image classes. The confidence vector is generated from the Softmax function and presents the model's prediction of the input. By concatenating the confidence vectors of all input images, we get the model's confidence matrix of size (M, N) , which includes the model's predictions of the entire input dataset of size M .

Based on the confidence matrix, the distance matrix $distMat$ of the N image classes is generated as described in Algorithm 1. Firstly, the distance matrix $distMat$ of the N image classes is initialized as a zero matrix of size (N, N) . Then we assign the class of each input image to $curClass$ and each image's confidence matrix of size $(1, N)$ to P (line 4-5). After that, we iterate over the $distMat$ and update the value through the iteration among P (line 8-14). Then, we calculate the $distMat$ using the iteration results and the iteration counts $distMatCount$ (line 16-17). Finally, we apply dimensionality reduction to the resulting matrix $distMat$ using t-SNE [45] to generate the 2D projection matrix for the distribution graph.

The distribution graph presents the distribution of the N image classes with respect to the predictions of each CNN model. Since the ImageNet structure is based on the WordNet hierarchy, there are eight root classes representing how human beings classify the

Algorithm 1 Constructing Distance Matrix of N ImageNet Classes

Input:

The image class list of all images in the dataset, $imgClasses$;
 The confidence matrix of the model, $confMat$;

Output:

The distance matrix of N image classes, $distMat$;

```

1:  $distMat \leftarrow zeros((N,N))$ 
2:  $distMatCount \leftarrow zeros((N,N))$ 
3: // iterate through all images
4: for  $imgIdx$  in  $range(M)$  do
5:    $curClass \leftarrow imgClasses[imgIdx]$  // get ground-truth class
6:    $P \leftarrow confMat[imgIdx, :]$  // get confidence score vector
7:   // iterate through all classes
8:   for  $compClass$  in  $range(N)$  do
9:      $confScore \leftarrow P[compClass]$ 
10:     $distMatCount[curClass, compClass] += 1$ 
11:    // calculate distance
12:    if  $compClass \neq curClass$  then
13:       $distMat[curClass, compClass] += (1 - confScore)$ 
14:    end if
15:  end for
16: end for
17:  $distMat \leftarrow divide(distMat, distMatCount)$  // get the average
18:  $distMat \leftarrow [distMat + transpose(distMat)]/2$ 
    
```

$N = 1000$ image classes in the ImageNet dataset. In the distribution graph of our system, each root class is represented by a specific color, allowing users to compare the model’s classification and human’s classification easily.

4.3 Visual Explanation Methods

As discussed in Sec. 2.3, visual explanation methods, especially the class-discriminative ones, can help novice ML practitioners to understand CNN model behaviors. Because they can highlight specific regions on the input image that is inferred to contribute the most to the model’s decision-making. However, most of the existing visual analytics tools for CNN comparison do not include any

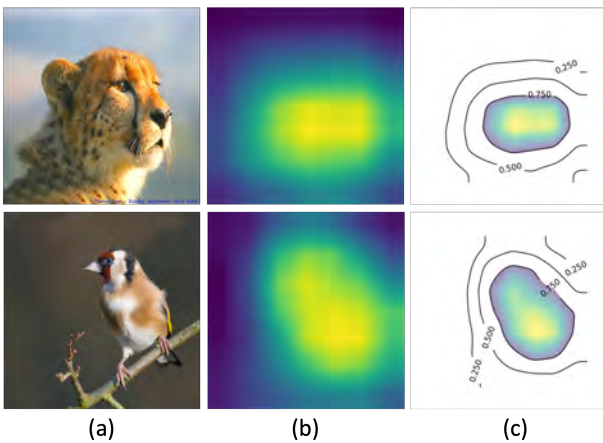


Fig. 3. Examples of how we improve the presentation of the visual explanation method’s results in VAC-CNN. Images at the same row are associate with the same original image. From left to right: (a) The original images; (b) The conventional visualizations of a visual explanation method; (c) The improved visualizations in VAC-CNN.

visual explanation methods. To fill this gap, we include five class-discriminative visual explanation methods in VAC-CNN, including Grad-CAM [17], BBMP [18], Grad-CAM++ [19], Smooth Grad-CAM++ [20], and Score-CAM [21]. Examples of these five methods are shown in Fig. 1 (D). The reason why these five methods are included is to cover multiple kinds of methods such as gradient-based explanations (Grad-CAM, Grad-CAM++, Smooth Grad-CAM++), perturbation-based explanation (BBMP), and score-based explanation (Score-CAM), which supports our design goal **G2**. Our analytics system is designed to be extensible, so other visual explanation methods can be easily added.

To achieve our design goal **G2**, we consolidate the presentation of the visual explanation method’s result. As shown in Fig. 3 (b), a conventional approach to present the visual explanation method’s results is showing the heatmap, which doesn’t provide any direct quantitative information. Thus, the subtle difference among multiple heatmaps can be hard to identify, making it not informative enough for the model comparison task.

In VAC-CNN, we add the quantitative information about the visual explanation method’s result by overlaying multiple contour lines over the heatmaps [46], which are associated with the attention matrix generated by the visual explanation method (with attention scores of $[0, 1]$, 0 for “no attention”). To support the highlighting of ROI, we also add a customizable threshold for users to remove regions of little attention accordingly. For example, a threshold of 0.5 means the region with attention scores lower than 0.5 will not be highlighted. As shown in Fig. 3(c), our improved visualizations of the explanation results incorporate qualitative information and quantitative measures of the attention level, which can support users in model comparison tasks more effectively.

4.4 Similarity Matrix Generation

When comparing multiple models based on a single image, users can benefit from a similarity matrix that intuitively shows the correlation of visual explanation methods’ results for the CNN models. We demonstrate the method to construct such similarity matrix in Algorithm 2.

In this algorithm, the generated saliency map from visual explanation methods are stored as matrices in a list $expResults$. We provide multiple widely-used image similarity measurements, including the structural similarity index (SSIM), the mean-square error (MSE), the $L1$ measure, and the hash function. The default

Algorithm 2 Constructing Similarity Matrix of Selected Models

Input:

The list of visual explanation results of models, $expResults$;
 The function of computing similarity scores, $simFunc$;

Output:

The similarity matrix of selected models, $simMatrix$;

```

1:  $L \leftarrow len(expResults)$ 
2:  $simMatrix \leftarrow zeros((L,L))$ 
3: for  $idx1$  in  $range(len(expResults))$  do
4:   for  $idx2$  in  $range(len(expResults))$  do
5:     // iterate through all visual explanation results
6:      $expRes1 \leftarrow expResults[idx1]$ 
7:      $expRes2 \leftarrow expResults[idx2]$ 
8:      $simMatrix[idx1, idx2] \leftarrow simFunc(expRes1, expRes2)$ 
9:   end for
10: end for
    
```

similarity measurement is set to $L1$ because of its wide acceptance, and other options are provided for users to select a different rule as needed. Based on the user-specified similarity comparison rules, we use the corresponding function *simFunc* to calculate the similarity score of two visual explanation results. After iterating over every element of *expResults*, we can get the similarity matrix *simMatrix* quantifying the similarity of each pair of the visual explanation results. To represent the value intuitively, we use *seaborn* [47] to generate the resulting matrix’s heat map. Then, users can interactively compare the behaviors of the selected CNN models through our designed interface described in Sec. 5.

4.5 Image Statistical Analysis

In some circumstances, the conventional visual explanation method may fail to provide enough information to explain the model’s prediction. For example, when the prediction result is wrong but the localization is correct, visual explanation methods doesn’t explain why the model made a wrong decision. To solve this problem, we go one step further by analyzing the information generated from the image region that the model cites as essential.

As discussed in [48], [49], CNN classifiers pre-trained on ImageNet have been proved to rely on texture information rather than the global object shape. However, current algorithms using image texture are often deep-learning-based [50], [51], which can severely interfere our system’s response speed. In VAC-CNN, we apply color intensity histograms (CIHs) to measure image information, which are commonly used to analyze the image content and evaluate the image similarity [52], [53], [54]. In this way, the analysis results can be generated in real-time (G4). Our process of image statistical analysis is shown in Fig. 4. Based on the model’s explanation, which highlights a specific image region essential for the model to make predictions, we can filter the original image by removing the “inessential” part. Then we visualize color intensity information of the filtered image (Fig. 4 (C)) to depict the statistical details of the image region that the model cites as essential in making predictions.

As the supplementary information of the visual explanation result, the color intensity histogram can help users further analyze what information the model extracts from the input. Through comparing the visual explanation results and the color intensity histograms, users can gain more insights into the underlying behaviors of the deep CNN model.

5 SYSTEM INTERFACE

To achieve our design goals described in Sec. 3, we integrate the techniques introduced in Sec. 4 into a web-based visual




Fig. 4. Example of the image statistical analysis process. From left to right: (a) A visual explanation method highlights ROI on the sample image; (b) The filtered image by only keeping the highlighted region of (a); (c) The color intensity histogram of the filtered image in (b).

analytics system, VAC-CNN, for the comparative studies of deep CNN models. As shown in Fig. 1, the system interface includes five primary views: “Overall Information View” (A), “Distribution Graph View” (B), “Task Selection Sidebar” (C), “Visual Explanation View” (D), and “Supplemental View” (E). In this section, we illustrate how these views coordinate to facilitate the three phases of comparison workflow described in Sec. 4.1.

5.1 Information Overview

In order to assist non-experts ML practitioners (G1), VAC-CNN provides an information overview for users to explore high-level CNN model performances and the general behaviors of multiple visual explanation methods. The analysis in this phase requires information from View (A), (B), (D), and (E) in our visual interface.

The Overall Information View (Fig. 1 (A)) illustrates the overall and detailed quantitative information of the included CNN models with multiple visualizations. The scatterplot labeled as (A1) indicates each model’s complexity and overall accuracy on the entire ImageNet validation set, where each point represents a CNN model. The radar chart labeled as (A2) reveals the accuracy performance of the selected models on the eight root classes. Each line of the radar chart represents one model’s performance, and the selectable legend located at the right of the chart enables users to remove uninterested models and only compare selected ones. Additionally, our interactive design allows users to change the pinned model or the pinned root class by a simple click, which can update the two zoomable bar charts shown at (A3) and (A4) of Fig. 1, representing leaf classes’ accuracy information of the model and root class, respectively, where the leaf classes are ranked in descending order of their accuracies. Thus, each part of the Overall Information View can work synergistically to illustrate each CNN model’s quantitative information from multiple aspects, helping users perceive models’ performances and achieve efficient high-level multi-model screening accordingly.

The Distribution Graph View (Fig. 1 (B)) reveals the distribution of the 1000 ImageNet classes. Each point represents a single image class, and the colors correspond to eight root classes. Generated from each model’s confidence score matrix, this visualization presents the model’s class-level behavior, enabling users to discover the model’s coherent or inconsistent behaviors across clusters of image classes. Besides, by looking at the clusters, users can also discover typical image class groups for further investigation in the following phases, which means this view also serves as a class recommendation. Similarly, smooth user interactions, including hovering over, clicking, zooming, etc., are supported as well.

As discussed in Sec. 4.3, the Visual Explanation View (Fig. 1 (D)) presents the example results of multiple visual explanation methods, informing non-expert users how each of the visual explanation method’s result looks like.

Finally, the Supplemental View (Fig. 1 (E)) provides users with supplemental information. At the information overview phase, two bar charts are presented at this view before users make any ImageNet class selections at the Task Selection Sidebar (Fig. 1 (C)). The first bar chart, “Range of Class Accuracy”, visualizes the range of the thirteen models’ accuracies on six image classes, including image classes on which the models have either diverging or parallel performances. And the second bar chart, “Average of Class Accuracy”, includes information related to six image classes, on which the models have coherent good performances or bad

performances. These two bar charts illustrate image classes with abnormal statistical characteristics, suggesting interesting images for users to explore in more detail.

5.2 Task Customization

VAC-CNN also supports users to customize the comparative study (G3) with the Task Selection Sidebar at the bottom left of our system interface (Fig. 1 (C)). From this view, users can select multiple CNN model(s), ImageNet class(es), visual explanation method(s), etc. Based on different selections, multiple subtasks can be performed in the following phase, including comparing multiple models over a particular image class, investigating a single model’s behaviors on multiple image classes, and explaining a single model’s behavior on images within a particular class. For the multi-model comparison task, VAC-CNN supports the users to select up to 13 models for comparison.

5.3 Model Investigation & Comparison

In the model investigation & comparison phase, the views (D) and (E) will be updated to present information based on the result from the user-specified comparison task (G3).

In the Visual Explanation View, various information is presented through a table format representation to better achieve the design goal G3. With multiple rows, this table presents comparison results of up to 13 models selected by users through task customization (as described in Sec. 5.2). Besides, the interaction features allow users to sort on multiple quantitative columns and search specific information to filter the results and get a deeper understanding.

We present an example to demonstrate what information is presented in this table. For instance, in the single-model investigation task described in Sec. 6.2, the view presents information including:

- the quantitative performance measures, such as model’s overall accuracy, class accuracy, confidence score, etc.;
- the corresponding information useful for understanding and comparing, such as the model name, image’s ground-truth class, and predicted class;
- the visual explanation results presented as contour plots, explanations on original images, as well as the CIH for the highlighted image region, etc.

Specifically, the CIH is used for supporting the single-model investigation tasks, so VAC-CNN only presents CIH when users are investigating a single model, as shown in Fig. 6 (D). As discussed in Sec. 4.3, VAC-CNN enables threshold adjustment for users to update the threshold of contour visualization of the visual explanation results. VAC-CNN coordinates the above information to support the comparative study process (G2).

In the Supplemental View, users can find various supplementary information to support model comparison and investigation according to different analysis needs. When users compare multiple models, this view includes information such as the original image selected by users, the similarity matrix of the models’ visual explanation results, and the scatterplots presenting the models’ accuracies on each selected image class (Fig. 5 (2-E)). When users investigate a single model, this view only shows the accuracy scatterplots of each selected image class since most of the essential information is already available in the Visual Explanation View.

6 USE CASES

In this section, we demonstrate how VAC-CNN can help novice ML practitioners conduct comparative studies with two use cases: (1) comparing the behavior of multiple models on the same image, (2) investigating a single model’s behavior on different images. The first use case demonstrates how VAC-CNN supports multi-step model comparison, from high-level screening over 13 models to the in-depth interpretable comparison of 7 models. The second use case is about single-model inspection, showing how our provided informative visual explanation assists users.

6.1 Comparing Multiple Models on Same Image

Alice is a Master’s student majoring in animal behavior study. She gained some basic knowledge about CNN and deep learning from a public course provided by the Computer Science Department and wants to apply it to her own major. Therefore, she uses VAC-CNN to explore the performances and behaviors of multiple CNN models on a group of images about animals.

After opening up the system, Alice starts by deciding the models, ImageNet classes, and visual explanation methods for her comparison task. She looks over the different plots in the Overall Information View (Fig. 1 (A)) to inspect the performances of the 13 CNNs and becomes interested in ResNet models when she notices the performance boost from *resnet18* to *resnet152*. She also notices that ResNet architectures often have good performance on the “animal” group from the radar chart in this view. Moreover, Alice finds “animal” forms better cluster for *resnet50* from the Distribution Graph View (Fig. 1 (B)), so she adds *resnet50*, *resnet101*, *resnet152* into the list of models. Then she looks at the first bar chart in the Supplemental View (Fig. 1 (E)) and finds that the models’ accuracies vary significantly on class “124 *crayfish*”, which belongs to the “animal” group, so she decides to choose this class for model comparison. Finally, Alice explores the Visual Explanation View (Fig. 1 (D)) and notices the ROI provided by “Grad-CAM” is very clear in general, so she decides to use “Grad-CAM” as the visual explanation method.

After having the models, ImageNet classes, and visual explanation methods she wants to select in mind, Alice moves on to the Task Selection Sidebar to customize her comparison task (Fig. 1 (C)). When restricting the ImageNet class selection to “124 *crayfish*”, Alice notices that a scatter plot in the Supplemental View is updated, as shown in Fig. 1 (E). And one model with remarkably bad performance, *alexnet* (14%), stands out. Besides, there are 3 other models whose accuracy is lower than 50%, *shufflenet_v2_x0_5* (32%), *squeezenet1_1* (36%), and *mobilenet_v2* (48%). Curious about the reasons behind those models’ failures, Alice also decides to add them to the model list for comparison. In this way, Alice has finalized the objectives of the model comparison task with 7 models:

- Models: *resnet50*, *resnet101*, *resnet152*, *alexnet*, *shufflenet_v2_x0_5*, *squeezenet1_1*, *mobilenet_v2*;
- ImageNet Class: 124 *crayfish*;
- Visual Explanation Method: Grad-CAM.

After finishing all of the customizations, Alice starts the comparison with the Visual Explanation View (Fig. 1 (D)) and the Supplemental View (Fig. 1 (E)). She first looks over the original images within the selected class (Fig. 5 (1-D)) from the Visual Explanation View. She finds that the image background of the main object—“*crayfish*”—is very complicated for almost every image in this class, which can be a possible cause of the varied

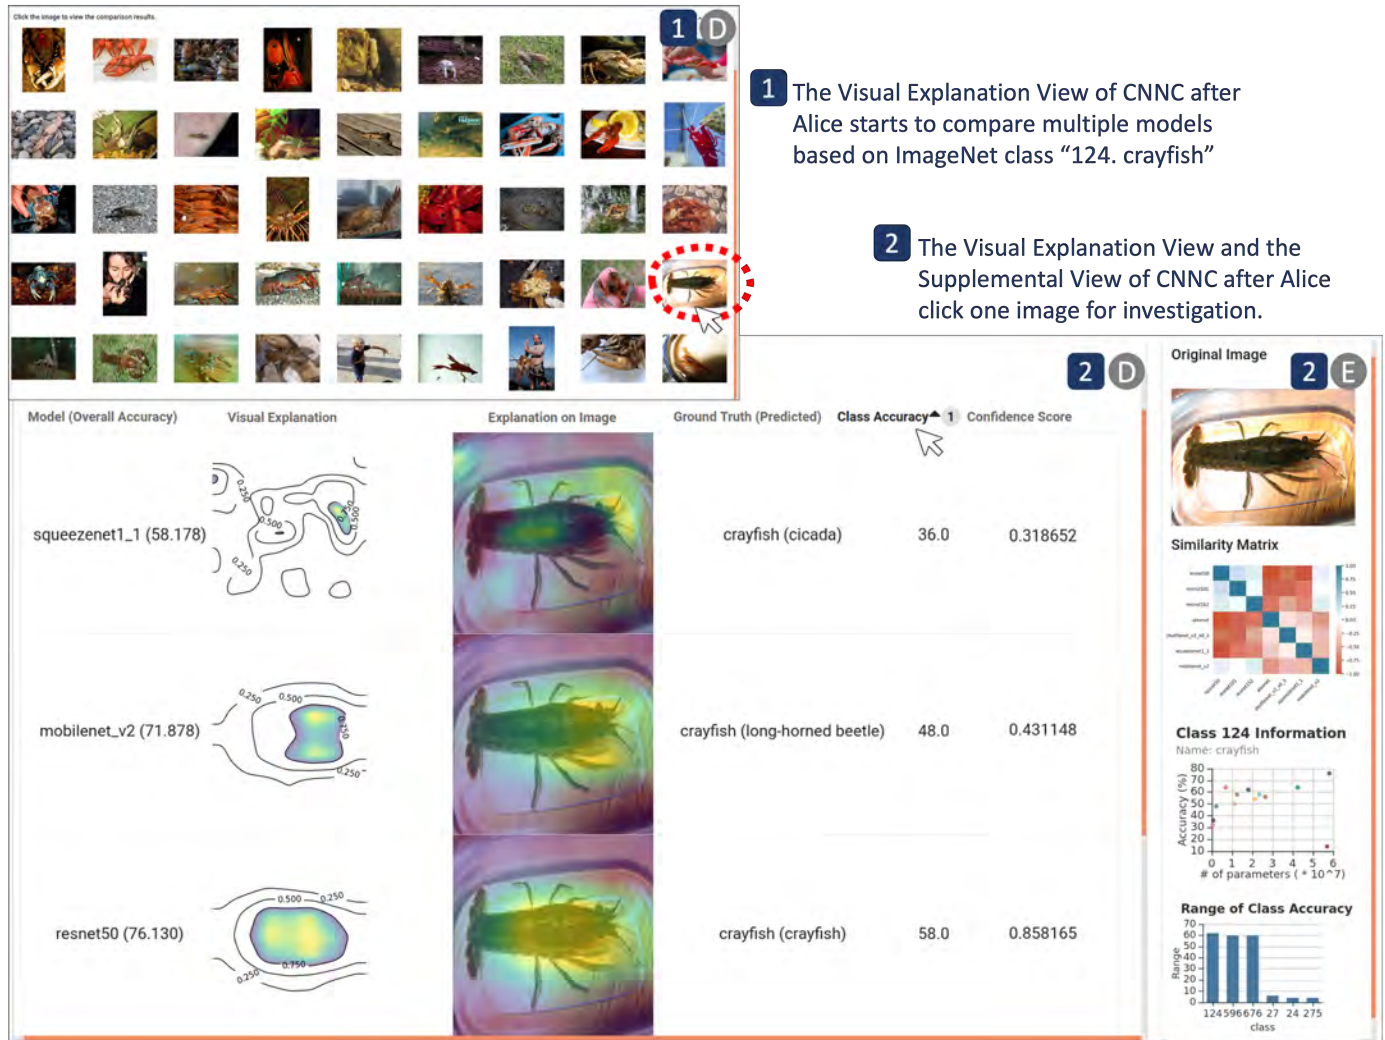


Fig. 5. Comparing the performance of 7 models on a selected image from class “124 *crayfish*” using VAC-CNN (Sec. 6.1). Looking over in (1-D), we can find a common characteristic of images in this class is their complicated background. After selecting one image of a shrimp, we can find more information with the updated views (2-D) and (2-E) (The views are marked as (D) and (E) to keep consistent with Fig. 1). By sorting the result in the table of (2-D) with different metrics, we can check the potential relationship between the model’s prediction performances and the ROI size. (Note: We compare 7 models in this task, but only show 3 models’ result in (2-D) because of space limitation.

model performances. With this hunch, Alice clicks on one image and begins to compare the models’ behaviors with the updated Visual Explanation View and Supplemental View (Fig. 5 (2)). As shown in Fig. 5 (2-D), by sorting the table according to the class accuracy, Alice inspects the visual explanation method’s results and the associated numerical information of the 7 models. She notices that, for the 3 models shown in Fig. 5 (2-D), *resnet50* is the only model that correctly classifies the input, while both *squeezenet1_1* and *mobilenet_v2* make incorrect predictions. By inspecting the visual explanation methods’ results, Alice realizes that the size of each model’s ROI has a positive relationship to the model’s prediction correctness: *alexnet* (lowest class accuracy) only highlights a very small region while the ROI of *resnet152* (highest class accuracy) is among the largest ones. After checking more images, Alice confirms the consistency of this observation. Given most of the images in this class have complicated backgrounds, Alice concludes that models with smaller views (i.e. smaller ROI) can’t perform very well in this object classification task. From this comparative study, Alice learns that when the images she is dealing with have complicated backgrounds, she will consider selecting

CNNs with broader views (e.g., *resnet50*) over others.

6.2 Investigating a Single Model on Different Images

This use case involves Bob, a first-year Ph.D. student majoring in Computer Science. He is developing a bird recognition App for a course project. And he wants to find the best model for the bird image classification function in his App.

Similar to Alice, Bob starts by deciding the model, ImageNet classes, and visual explanation methods for his task. He first checks the models’ differences in complexities and overall accuracies with the scatter plot in the Overall Information View (Fig. 6 (A1)). He finds that *resnet152* achieves the best performances compared to other CNN models. And such an advantage is particularly prominent with the root classes “animal” and “fungus” according to the radar chart in the Overall Information View (Fig. 6 (A2)). Therefore, Bob decides to choose *resnet152* as the model to dig deeper into. As shown in Fig. 6 (B), then he zooms into the distribution graph of *resnet152* to check the cluster of “bird” species and decides to select class “130 *Flamingo*” to explore model behaviors on it. Finally, after looking over the examples

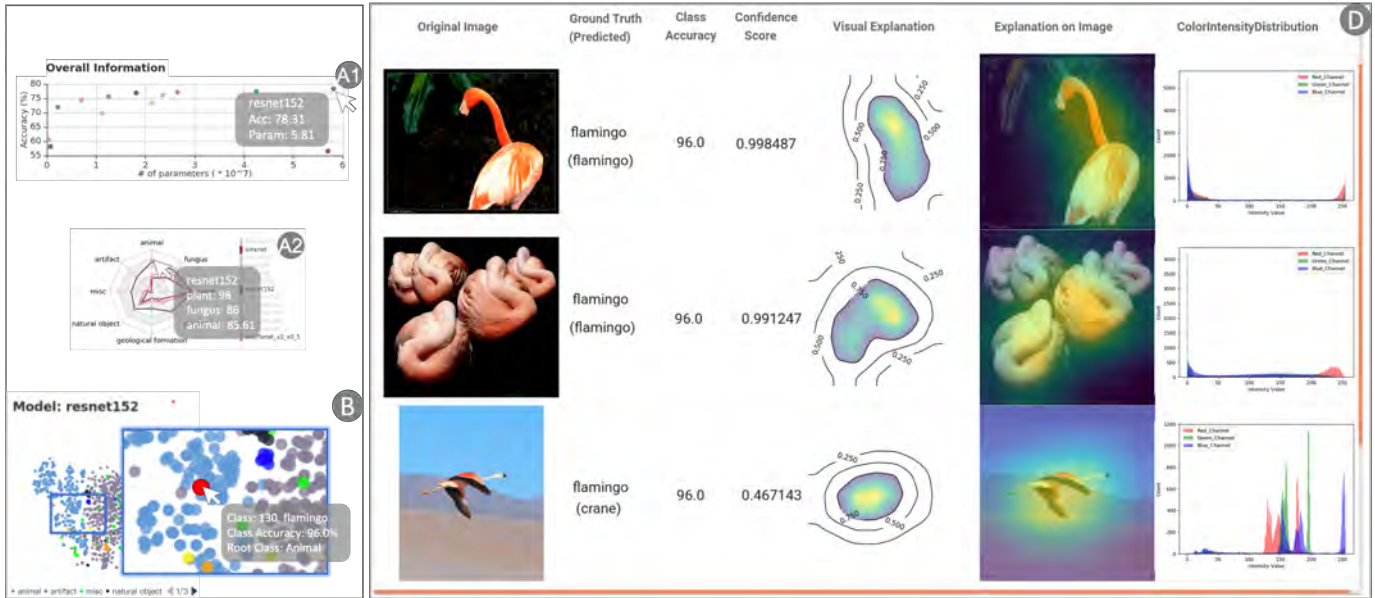


Fig. 6. Investigating a single model *resnet152* on different images of class “flamingo” using VAC-CNN (Sec. 6.2). The multiple visualizations in (A1), (A2), and (B) can help users notice the high-level behavior pattern of *resnet152* from different perspectives. And the coordinated view (D) shows the CIH results to help users discover the common patterns among models’ failure cases on the “flamingo” class. (The views are marked as (A1), (A2), (B), (D) to keep consistent with Fig. 1.)

of multiple visual explanation methods in the Visual Explanation View (Fig. 1 (D)), he chooses “Smooth Grad-CAM++” as the visual explanation method for model interpretation. In summary, the objectives of Bob’s model comparison task are:

- Model: *resnet152*;
- ImageNet Class: 130 *Flamingo*;
- Visual Explanation Method: Smooth Grad-CAM++.

After customizing his comparison task with the Task Selection Sidebar, Bob starts the model investigation from the updated view (Fig. 6 (D)). He first notices that Smooth Grad-CAM++ indicates correct localization of the main object in every image in the class “130 *Flamingo*”, even for those incorrectly predicted ones. He feels excited about this discovery and continues to look for the cause of those incorrect predictions made by *resnet152*. He finds that *resnet152* correctly classifies the first two images with high confidence scores but misclassifies the third as “Crane” in Fig. 6 (D). In contrast, Bobs thinks the second image is more challenging to recognize than the third one in his eyes. He tries to explain this phenomenon from the color intensity histograms (CIH) provided by VAC-CNN. By comparing the CIHs of the three images, he realizes that the second image’s CIH is highly similar to the first one, while the third image’s CIH looks much more different from the other two (Fig. 6 (D)). After checking the conditions with other image classes of bird species, Bob finds such observation still holds for most failure cases. He shares this interesting discovery with his course instructor. The instructor suggests he construct a small subgroup of the bird classes that most confuses *resnet152*, apply data augmentation specifically, and use it to fine-tune the model. Bob optimizes his model following this idea, and makes his bird recognition App more potent in the classification task.

7 PRELIMINARY EVALUATION STUDY

VAC-CNN is designed for assisting novice ML practitioners in comparing and understanding multiple CNN models. In this section,

we conducted a preliminary evaluation study to demonstrate the usefulness of our system. Specifically, we intended to understand whether VAC-CNN was effective in helping users: (1) gain high-level understanding of various CNN models (G1); (2) interpret CNN behaviors (G2); (3) customize different comparison tasks (G3). We also investigated them about how they felt about the smoothness of the system as well as the interactions (G4). The evaluation of our study mainly adopts qualitative analysis towards participants’ behaviors and feedbacks, along with minor quantitative analysis of their self-reported ML knowledge level and rating scores of the system.

Considering the unprecedented challenging situation brought by Covid-19, our study environment was restricted and we had to do everything remotely with limited number of participants. However, because we carefully designed the entire study procedure and address a thorough evaluation, the validity of VAC-CNN can still be proved through this study.

7.1 Participants

We recruited 12 participants (6 male, 6 female), including 7 M.S. students and 5 Ph.D. students. We asked them to self-report their familiarity with three areas on a scale of [0, 10] (0 for “no knowledge” and 10 for “expert”) and report the statistics as follows:

- Basic machine learning techniques: $Md = 4.00, IQR = 2.25$;
- CNNs: $Md = 2.50, IQR = 1.25$;
- Visual explanation methods: $Md = 2.00, IQR = 1.25$.

The result shows that all of the participants have limited deep learning and XAI background, so they belong to our target user group, novice ML practitioners.

7.2 Task Design

We asked each participants to perform the same tasks using VAC-CNN and observe their behavior patterns during the process. After

getting familiar with the visual interface, they were asked to perform the following tasks:

T1 Browse high-level information: The participants were asked to get a high-level understanding of model performances and the behaviors of multiple visual explanation methods (**G1**) through interactions with multiple visualizations presented in our visual interface (**G4**). They were encouraged to use as much interactions as possible and describe their findings.

T2 Compare multiple models: To observe how VAC-CNN can assist users in multi-model comparison, we asked the participants to compare at least two models (**G2, G3**). The models, as well as other customizable options, such as visual explanation methods, were chosen by the participants. And we asked them to provide the reason of their selections. The participants were asked to identify common and unique behaviors of the compared models, and which components of VAC-CNN lead to their findings.

T3 Investigate a single model: In this task, the participants were asked to select one CNN model for in-depth investigation. Similar with task T2, we asked them to decide all customizable options, including the model they chose to investigate, and provide us with the reasons (**G3**). The participants were asked to describe their understanding of model behaviors and how VAC-CNN assist them during the process (**G2, G4**).

7.3 Study Setup and Procedure

Our preliminary evaluation study is conducted remotely through one-on-one video meeting with each participant. The participants were asked to access VAC-CNN running at a remote server with their personal computers. Before the study started, we asked each participant to self-report their knowledge background and basic demographic information. At the beginning of the study, we provided a 5-minute tutorial session to introduce the models, dataset, visual components, and interactions built in VAC-CNN. After that, we asked the participants to perform the three tasks described in Sect. 7.2, and encourage them to use as many system components as possible. This session took around 30 minutes on average and participants followed the think-aloud protocol when they performed these tasks. Finally, the participants were invited to fill up a usability questionnaire and share their feedback about experiences with VAC-CNN in a 5-minute follow-up interview.

7.4 Results and Discussion

This section demonstrates our findings from the usability questionnaire, the follow-up interview, and the behavior observation of all users. We asked the participants to rate the usability of the system in the questionnaire as well as collected their comments about the system in a follow-up interview. The result shows that we successfully achieved all of our design goals, but it also reveals some shortcomings that can be improved in the future.

The questionnaire includes two quantitative questions: rating the easy-to-use level and the helpful level of our system. When rating how our system is easy to use on a scale of [0, 10] (0 for “very difficult”, 10 for “very easy”), the participants provided scores with $Md = 8$, $IQR = 2.25$, and more than 60% of our participants’ rates are 8 or higher. When rating the helpful level of our system on a scale of [0, 10] (0 for “absolutely not helpful”, 10 for “absolutely helpful”), the participants provided scores with $Md = 6$, $IQR = 1.5$, and more than 75% of our participants’ rates are 6 or higher.

Our observation of the user behavior and the comments we received from the interview show that most of our design goals are fulfilled well. All participants were able to finish task T1, which means they can generate high-level insights of models, image classes, and visual explanation methods through exploring VAC-CNN (**G1**). One common behavior pattern of the participants was using the sortable table to investigate visual explanation results and the corresponding numerical information, through which they interpret model behaviors and answered our questions at tasks T2 and T3 (**G2, G3**). Most of the participants (9 out of 12) mentioned that they enjoyed the smooth interface, and 4 of them thought the real-time presents of the visual explanation results were impressive (**G4**). “*I like the way how multiple views are coordinated. I can start investigate a new model through a simple click*”, commented by participant P4.

However, the results also reflect some shortcomings of our system. A few participants (2 out of 12) only had limited interactions with the distribution graph view, because they were not familiar with clustering and felt it was hard to identify model behaviors through this visualization. Participant P9 felt “*understanding a model’s behavior pattern from this (view) is hard for me*”. Some of the participants (3 out of 12) mentioned in the interview that the CIH might not provide convincing results in some scenarios, and one participant thought the system can be improved by including collective analysis towards visual explanation methods on the entire dataset. We will discuss how to address these problems in Sec. 8.

8 LIMITATIONS AND FUTURE WORK

Through our preliminary evaluation study, we identify a few limitations of VAC-CNN. In this section, we discuss these limitations and the corresponding future work.

Image statistical analysis. The image statistical analysis functionality is supposed to support model behavior comparison when visual analytics methods fail. However, we have found that there are many conditions when the Color Intensity Histograms can not provide convincing supplemental information for understanding model behaviors. In the future, we plan to experiment with a new approach to image texture analysis in real-time, which should be robust and effective in various application scenarios.

Collective model evaluation. Our current system includes thirteen CNN models and five visual explanation methods. Although we support customized comparing tasks on multiple CNN models, we don’t provide a collective model evaluation. In the future, we plan to extend our work by introducing model behavior evaluation on the dataset level, with which users are able to obtain a high-level evaluation of model behaviors across the entire dataset as well as explore specific behaviors on single instances.

Precise evaluation of qualitative comparisons. Our system assists researchers in combining both quantitative and qualitative analysis and allows users to update results interactively. However, despite adding contour visualization to quantify visual explanation results, judging behavior differences of models is still largely observation-based, which could be imprecise. In the future, we plan to incorporate quantitative measures to support evaluation, such as showing the amount of noise in the visual explanation outputs or the accuracy of the highlighted region.

Customization Recommendation. To support interpretable CNN model comparisons, our system includes multiple class-discriminative visual explanation methods and presents examples on each of them. Although customizable options can support

insight-building by providing various tryouts, our system would be more user-friendly (ML novices in particular) if it could recommend explanation methods according to users' demands. As future work, we plan to design recommendation strategies, such as building evaluation matrices of the visual explanation methods according to the data randomization test [55], to assist ML novices in choosing visual explanation methods.

9 CONCLUSION

In this paper, we present a visual analytics system VAC-CNN (Visual Analytics for Comparing CNNs) to assist novice ML practitioners in the comparative studies of deep Convolutional Neural Networks. To support model interpretability, VAC-CNN integrates multiple visual explanation methods and improves the result visualization. The system coordinates quantitative measures and informative visual explanations, and supports flexible customization of the model exploring tasks, including multi-model comparison and single-model investigation. We evaluate the usability of VAC-CNN in supporting ML beginners through a preliminary evaluation study. We hope our work will encourage further exploration of the inner behaviors of CNN models, and inspire the design of the next generation CNN comparison tools.

ACKNOWLEDGMENTS

This research is supported in part by the U.S. National Science Foundation through grant IIS-1741536 and a gift grant from Bosch Research. We would like to thank all the participants of our preliminary evaluation study during this challenging time. We also want show our gratitude to Norma Gowans for narrating in our demonstration video. We appreciated Takanori Fujiwara, Jianping (Kelvin) Li, and Qi Wu for their precious suggestions that improve this work. We wish to extend our special thanks to anonymous reviewers for their thoughtful feedbacks and comments.

REFERENCES

- [1] D. Bhatt, C. Patel, H. Talsania, J. Patel, R. Vaghela, S. Pandya, K. Modi, and H. Ghayvat, "CNN variants for computer vision: History, architecture, application, challenges and future scope," *Electronics*, vol. 10, no. 20, p. 2470, 2021.
- [2] Y. LeCun, L. Jackel, L. Bottou, C. Cortes, J. S. Denker, H. Drucker, I. Guyon, U. A. Muller, E. Sackinger, P. Simard *et al.*, "Learning algorithms for classification: A comparison on handwritten digit recognition," *Neural networks: the statistical mechanics perspective*, vol. 261, p. 276, 1995.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [4] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [5] F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, "SqueezeNet: Alexnet-level accuracy with 50x fewer parameters and <0.5 mb model size," *arXiv preprint arXiv:1602.07360*, 2016.
- [6] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.
- [7] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.
- [8] X. Zhang, X. Zhou, M. Lin, and J. Sun, "Shufflenet: An extremely efficient convolutional neural network for mobile devices," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6848–6856.
- [9] Y. Hu, G. Wen, M. Luo, D. Dai, J. Ma, and Z. Yu, "Competitive inner-imaging squeeze and excitation for residual network," *arXiv preprint arXiv:1807.08920*, 2018.
- [10] M. F. Aydogdu, V. Celik, and M. F. Demirci, "Comparison of three different CNN architectures for age classification," in *2017 IEEE 11th International Conference on Semantic Computing (ICSC)*. IEEE, 2017, pp. 372–377.
- [11] W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu, and F. E. Alsaadi, "A survey of deep neural network architectures and their applications," *Neurocomputing*, vol. 234, pp. 11–26, 2017.
- [12] A. Khan, A. Sohail, U. Zahoor, and A. S. Qureshi, "A survey of the recent architectures of deep convolutional neural networks," *arXiv preprint arXiv:1901.06032*, 2019.
- [13] A. Canziani, A. Paszke, and E. Culurciello, "An analysis of deep neural network models for practical applications," *arXiv preprint arXiv:1605.07678*, 2016.
- [14] A. Mukhopadhyay, P. Biswas, A. Agarwal, and I. Mukherjee, "Performance comparison of different cnn models for indian road dataset," in *Proceedings of the 2019 3rd International Conference on Graphics and Signal Processing*, 2019, pp. 29–33.
- [15] D. Gunning, "Explainable artificial intelligence (xai)," *Defense Advanced Research Projects Agency (DARPA), nd Web*, vol. 2, 2017.
- [16] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2921–2929.
- [17] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [18] R. C. Fong and A. Vedaldi, "Interpretable explanations of black boxes by meaningful perturbation," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3429–3437.
- [19] A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian, "Grad-CAM++: Generalized gradient-based visual explanations for deep convolutional networks," in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 839–847.
- [20] D. Omeiza, S. Speakman, C. Cintas, and K. Weldermariam, "Smooth Grad-CAM++: An enhanced inference level visualization technique for deep convolutional neural network models," *arXiv preprint arXiv:1908.01224*, 2019.
- [21] H. Wang, M. Du, F. Yang, and Z. Zhang, "Score-CAM: Improved visual explanations via score-weighted class activation mapping," *arXiv preprint arXiv:1910.01279*, 2019.
- [22] J. Krause, A. Perer, and K. Ng, "Interacting with predictions: Visual inspection of black-box machine learning models," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2016, pp. 5686–5697.
- [23] H. Zeng, H. Haleem, X. Plantaz, N. Cao, and H. Qu, "CNNComparator: Comparative analytics of convolutional neural networks," *arXiv preprint arXiv:1710.05285*, 2017.
- [24] J. Zhang, Y. Wang, P. Molino, L. Li, and D. S. Ebert, "Manifold: A model-agnostic framework for interpretation and diagnosis of machine learning models," *IEEE transactions on visualization and computer graphics*, vol. 25, no. 1, pp. 364–373, 2018.
- [25] S. Das, D. Cashman, R. Chang, and A. Endert, "Beames: Interactive multimodel steering, selection, and inspection for regression tasks," *IEEE computer graphics and applications*, vol. 39, no. 5, pp. 20–32, 2019.
- [26] S. Murugesan, S. Malik, F. Du, E. Koh, and T. M. Lai, "Deepcompare: Visual and interactive comparison of deep learning model performance," *IEEE computer graphics and applications*, vol. 39, no. 5, pp. 47–59, 2019.
- [27] D. L. Arendt, N. Nur, Z. Huang, G. Fair, and W. Dou, "Parallel embeddings: a visualization technique for contrasting learned representations," in *Proceedings of the 25th International Conference on Intelligent User Interfaces*, 2020, pp. 259–274.
- [28] C. Olah, A. Mordvintsev, and L. Schubert, "Feature visualization," *Distill*, vol. 2, no. 11, p. e7, 2017.
- [29] A. Mordvintsev, C. Olah, and M. Tyka, "DeepDream—a code example for visualizing neural networks," *Google Research*, vol. 2, no. 5, 2015.
- [30] K. Simonyan, A. Vedaldi, and A. Zisserman, "Deep inside convolutional networks: Visualising image classification models and saliency maps," *arXiv preprint arXiv:1312.6034*, 2013.
- [31] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," *arXiv preprint arXiv:1412.6806*, 2014.

- [32] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *European conference on computer vision*. Springer, 2014, pp. 818–833.
- [33] H. Talebi and P. Milanfar, "Nima: Neural image assessment," *IEEE Transactions on Image Processing*, vol. 27, no. 8, pp. 3998–4011, 2018.
- [34] P. E. Rauber, S. G. Fadel, A. X. Falcao, and A. C. Telea, "Visualizing the hidden activity of artificial neural networks," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 1, pp. 101–110, 2016.
- [35] M. Liu, J. Shi, Z. Li, C. Li, J. Zhu, and S. Liu, "Towards better analysis of deep convolutional neural networks," *IEEE transactions on visualization and computer graphics*, vol. 23, no. 1, pp. 91–100, 2016.
- [36] K. Wongsuphasawat, D. Smilkov, J. Wexler, J. Wilson, D. Mane, D. Fritz, D. Krishnan, F. B. Viégas, and M. Wattenberg, "Visualizing dataflow graphs of deep learning models in tensorflow," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 1–12, 2017.
- [37] A. Bilal, A. Jourabloo, M. Ye, X. Liu, and L. Ren, "Do convolutional neural networks learn class hierarchy?" *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 152–162, 2017.
- [38] M. Kahng, P. Y. Andrews, A. Kalro, and D. H. P. Chau, "A cti v is: Visual exploration of industry-scale deep neural network models," *IEEE transactions on visualization and computer graphics*, vol. 24, no. 1, pp. 88–97, 2017.
- [39] F. Hohman, H. Park, C. Robinson, and D. H. P. Chau, "Summit: Scaling deep learning interpretability by visualizing activation and attribution summarizations," *IEEE transactions on visualization and computer graphics*, vol. 26, no. 1, pp. 1096–1106, 2019.
- [40] L. Schubert, M. Petrov, S. Carter, N. Cammarata, G. Goh, and C. Olah, "Openai microscope," Apr 2020, <https://microscope.openai.com/models>.
- [41] Y. Ma, A. Fan, J. He, A. R. Nelakurthi, and R. Maciejewski, "A visual analytics framework for explaining and diagnosing transfer learning processes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 2, pp. 1385–1395, 2020.
- [42] Y. Li, T. Fujiwara, Y. K. Choi, K. K. Kim, and K.-L. Ma, "A visual analytics system for multi-model comparison on clinical data predictions," *Visual Informatics*, vol. 4, no. 2, pp. 122–131, 2020.
- [43] Z. Li, F. Liu, W. Yang, S. Peng, and J. Zhou, "A survey of convolutional neural networks: analysis, applications, and prospects," *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [44] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "ImageNet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [45] L. v. d. Maaten and G. Hinton, "Visualizing data using t-sne," *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [46] J. D. Hunter, "Matplotlib: A 2d graphics environment," *Computing in science & engineering*, vol. 9, no. 03, pp. 90–95, 2007.
- [47] M. L. Waskom, "Seaborn: statistical data visualization," *Journal of Open Source Software*, vol. 6, no. 60, p. 3021, 2021.
- [48] N. Baker, H. Lu, G. Erlikhman, and P. J. Kellman, "Deep convolutional networks do not classify based on global object shape," *PLoS computational biology*, vol. 14, no. 12, p. e1006613, 2018.
- [49] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel, "ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness," *arXiv preprint arXiv:1811.12231*, 2018.
- [50] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2414–2423.
- [51] W. Cai and Z. Wei, "Diversity-generated image inpainting with style extraction," *arXiv preprint arXiv:1912.01834*, 2019.
- [52] K. Roy and J. Mukherjee, "Image similarity measure using color histogram, color coherence vector, and sobel method," *International Journal of Science and Research (IJSR)*, vol. 2, no. 1, pp. 538–543, 2013.
- [53] A. Malakar and J. Mukherjee, "Image clustering using color moments, histogram, edge and k-means clustering," *International Journal of Science and Research*, vol. 2, no. 1, pp. 532–537, 2013.
- [54] S. Phetnuam and T. Yingthawornsuk, "Classification of categorized kmutt-bkt's landscape images using rgb color feature," in *2018 14th International Conference on Signal-Image Technology & Internet-Based Systems (SITIS)*. IEEE, 2018, pp. 327–331.
- [55] J. Adebayo, J. Gilmer, M. Muelly, I. Goodfellow, M. Hardt, and B. Kim, "Sanity checks for saliency maps," *arXiv preprint arXiv:1810.03292*, 2018.



Xiwei Xuan is a Ph.D. student in computer science at the University of California, Davis. Before UC Davis, she received the M.S. degree in electrical engineering in 2020 from the Washington University in St. Louis and received the B.S. and M.S. degrees in microelectronics from the Harbin Institute of Technology. Her main research interests include visual analytics, machine learning, and explainable artificial intelligence.



Xiaoyu Zhang is a Ph.D. candidate in computer science at the University of California, Davis. She received her B.S. in digital media art from Xiamen University and her M.S. in computer science from Zhejiang University. Her research interest lies in visual analytics and information visualization. More specifically, she studies data analysis and visualization techniques to explore and exploit underlying knowledge, fact, or pattern from large text, tabular or ontological data.



Oh-Hyun Kwon received the PhD degree in computer science from the University of California, Davis in 2021. His PhD research focused on developing machine learning and immersive approaches to graph visualization. He is currently working in the industry in the areas of data visualization, visual analytics, and machine learning.



Kwan-Liu Ma is a distinguished professor of computer science at the University of California, Davis. His research is in the intersection of data visualization, computer graphics, human-computer interaction, and high performance computing. For his significant research accomplishments, Ma received several recognitions including being elected as IEEE Fellow in 2012, recipient of the IEEE VGTC Visualization Technical Achievement Award in 2013, and inducted to IEEE Visualization Academy in 2019.