Interactive Visualization of Deep Learning for 3D Brain Data Analysis

Huang Li[1], Shiaofen Fang[1], Joaquín Goñi[2], Andrew J. Saykin[3], Li Shen[4]

[1]: Department of Computer & Information Science
Purdue School of Science
Indiana University Purdue University Indianapolis
Email: shfang@iupui.edu

[2]: School of Industrial Engineering & Weldon School of Biomedical Engineering
Purdue University
Email: jgonicor@purdue.edu

[3]: Center for Neuroimaging
School of Medicine
Indiana University
Email: asaykin@iupui.edu

[4]: Department of Biostatistics, Epidemiology and Informatics
Perelman School of Medicine
University of Pennsylvania
Email: Li.Shen@pennmedicine.upenn.edu

***Abstract*— With multiple hidden layers and massive combinations of features and weights, deep learning models are hard to understand, and even more difficult to interact with. In this paper we describe a visual analytics platform to help with the understanding of and interaction with the deep learning process of human brain image data. A brain connectome network dataset is used to train a classifier for the diagnosis of Alzheimer's Disease (AD). 3D rendering of brain images is integrated into the interactive visualization process of a deep neural network to bring contextual information of the application to the analysis framework. A backpropagation algorithm is applied to track the image features that are captured by each node in the hidden layers. Our results demonstrate that interactive visualization can not only help the understanding of the deep learning process, but also provide a platform for domain experts to interact with and assist in the learning process, which can potentially enhance the interpretability and accuracy of the analysis.**

Keywords— visualization, deep learning, human-computer interaction, brain imaging.

1 Introduction

Recent advances in deep neural networks and other machine learning techniques have led to unprecedented breakthroughs in many big data analysis problems and applications. The intrinsic complexity in deep learning, introduced by multiple hidden layers and massive combinations of features and weights, makes the model hard to interpret. Therefore, in most scenarios, deep learning models are used as black boxes [1].

Visualization has been playing an increasingly important role in data mining and data analytics. This transformation of visualization from data viewing to becoming an integrated part of the analysis process led to the birth of the field of visual analytics [2]. In visual analytics, carefully designed visualization processes can effectively "decode" the insight of the data through visual transformations and interactive exploration. Many successful applications of visual analytics have been published in recent years, ranging from bioinformatics and medicine to engineering and social sciences.

As automatic methods, machine learning algorithms, particularly deep learning algorithms, provide the users very little information about how and why the algorithms work or fail. The underlying deep learning models are also designed primarily for the convenience of learning from data, but they are not easy for the users to understand or interact with. Interactive visualization can provide an effective mechanism to help the users understand and interact with the deep learning process [3]. This has two important potential benefits.

(1) Understanding. Understanding the deep learning process in the context of the application is particularly important for biomedical applications in which how the biological processes and pathways are involved is sometimes more important than the result itself. It is also difficult to improve the efficiency and performance of the algorithms without a clear understanding of how and why the different components work in deep learning algorithms with many layers and interconnected components.

(2) Knowledge Input. In the biomedical field, human knowledge and expertise is critical in producing biologically meaningful results. Therefore, human input during the deep learning process is sometimes desirable and even necessary to generate the best and the most meaningful result.

Our goal is to develop an interactive visualization framework in a deep learning environment to provide contextual associations with the application such that the user can observe and understand the evolution and performance of the internal layers and their feature selections. Furthermore, user will also be able to interact with the system through the visualization tools to provide feedback that may alter the learning process, for example changing the weight of an edge or adding/removing a node.

In order to better understand and effectively interact with the deep learning process, the user needs to be able to associate the deep learning model's intermediate layers and their nodes with application specific information to make contextual judgements. This requires the visualization system to provide not only a visual representation for the deep neural networks, but also the application specific visual information associated with the various components of the neural networks. In other words, the system needs to be able to trace from the intermediate nodes and edges back to the original

This is the author's manuscript of the article published in final edited form as:
Li, H., Fang, S., Goñi, J., Saykin, A. J., & Shen, L. (2021, October). Interactive Visualization of Deep Learning for 3D Brain Data Analysis. In 2021 IEEE 20th International Conference on Cognitive Informatics & Cognitive Computing (ICCI* CC) (pp. 85-91). IEEE. https://doi.org/10.1109/ICCICC53683.2021.9811312

input data to generate meaningful information in the context of the application. Such contextual associations can also help the user make decisions on whether there is a valid reason to alter the learning process, and if so how this can be achieved through the editing operations of the intermediate layers.

In this paper, we propose an interactive visualization approach to the understanding of and interaction with a deep learning algorithm for 3D brain images data. The goal is to train a classifier for the diagnosis of Alzheimer's Disease (AD). Although Alzheimer's Disease is currently not curable, early diagnosis is extremely important for patients' quality of life. Furthermore, interactive visualization of the deep learning process may reveal interesting and meaningful biological interpretations that could provide potential benefits for the development of a cure. Our primary contributions in this paper include:

(1) An effective interactive visualization framework that provides direct visual associations of the deep neural network structures with the anatomical structures of the brain. It also provides user editing operations to alter the internal layers to re-train the model (with a new initial condition) when necessary.

(2) A recursive backpropagation algorithm to compute the weights of the input features in each node in the hidden layers, which can lead to the anatomical or biological interpretation of the hidden layer nodes.

In the following, we will first, in Section 2, discuss previous work related to visualization techniques in machine learning and deep learning. We will provide some background information regarding the dataset and related brain application in Section 3. The interactive visualization framework for deep learning of brain data application will be described in Section 4. Some experimental results and implementation details will be given in Section 5. We conclude the paper with final remarks and future work in Section 6.

II. RELATED WORK

There have been a number of previous works on using visualization to help understand the machine learning processes. Neural Networks received the most attention due to its "black box" nature of the learning model and the complexity of its internal components. Multi-dimensional visualization techniques such as scatterplot matrix have been used to depict the relationships between different components of the neural networks [4][5]. Typically, a learned component is represented as a higher dimensional point. The 2D projections of these points in either principal component analysis (PCA) spaces or a multi-dimensional scaling (MDS) space can better reveal the relationships of these components that are not easily understood, such as clusters and outliers. Several techniques have applied graph visualization techniques to visualize the topological structures of the neural networks [6][7][8]. Visual attributes of the graph can be used to represent various properties of the neural network models and processes.

Several recent studies tackle specific challenges in the visualization of deep neural networks due to the large number of components, connections and layers. In [9]. Liu et al. developed a visual analytics system, CNNVis, that helps machine learning experts understand deep convolutional neural networks by clustering the layers and neurons. Edge bundling is also used to reduce visual clutter. Techniques have also been developed to visualize the response of a deep neural network to a specific input in a real-time dynamic fashion [10][11]. Observing the live activations that change in response to user input helps build valuable intuitions about how convnets work. The DQNViz system [12] provides a visual analytics environment for the understanding of a deep reinforcement learning model. GAN Lab system [13] uses visualization to help non-expert users to learn how a Deep Generative Model works.

Several studies focus on visualizing the features captured by deep neural network. Class Activation Mapping (CAM) [14] was proposed by Zhou et al as a method for identifying discriminative locations used by the convolutional layers in deep learning model without any fully-connected layer. By extending CAM with gradient, Selvaraju et al proposed Grad-CAM [15] which works with fully-connected layers. But these methods did not reveal how the nodes in the hidden layer capture features from those discriminative locations and aggregate them together in the feed-forward process, which is important for users to make sense of the deep learning model.

Visualization and visual analytics methods have been proposed for the performance analysis of machine learning algorithms in different applications [16][17][18]. Interactive methods have also been proposed to improve the performance of machine learning algorithms through feature selection and optimization of parameter settings. Some general discussions are given in [19] and [20]. In [21], a visual analytics system for machine learning support called Prospector is described. Prospector supports model interpretability and actionable insights, and provides diagnostic capabilities that communicate interactively how features affect the prediction. In [22], a multigraph visualization method is proposed to select better features through an interactive process for the classification of brain networks. Other performance improvement methods include training sample selection and classifier tuning [23] and model manipulations by user knowledge [24][25][26].

More generic visualization methods for general machine learning models have been developed in recent years. The *Manifold* system [27] provides a generic framework that does not rely on or access the internal logic of the model and solely observes the input and output. An ontology, VIS4ML, is proposed in [28] for VA-Assisted machine learning. In [29], A generic visualization method for machine learning model is proposed to help select the optimal set of sample input data.

Compared to the existing methods on visualization in machine learning and deep learning applications, our approach emphasizes building association relationships between hidden layer nodes and the brain image features (phenotypes) such that user can observe and interact directly with the complex anatomical structures of the brain during the deep learning process. This work is also a good example of how to integrate scientific visualization and information visualization techniques in a deep learning visual analytics platform.

III. BRAIN IMAGE DATA SET

Data used in the preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative

(ADNI) database (adni.loni.usc.edu). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early AD. For up-to-date information, see www.adni-info.org.

For the proposed interactive visualization study, we downloaded the structural MRI (sMRI) and diffusion tensor imaging (DTI) scans as well as relevant clinical information from the ADNI website. Our application is the human brain network data analysis for Alzheimer's Disease classification. The DTI data was first denoised and corrected for motion and distortion using the approach described in a previous study [30]. Tractography was performed in Camino [31] based on white matter fibre orientation distribution function (ODF). Streamlines were modelled with a multi-tensor modelling approach, where voxels will fit up to two fibre orientations. Second, the sMRI scans were registered to the b0 volume of DTI data using the FNIRT toolbox [32] and 278 brain regions of interest (ROIs) were extracted following [33]. The final networks between 278 brain ROIs were constructed using fibres going through white matter and connecting ROIs.

In this project, we focus on the link level measures and calculated the fibre anisotropy (FA), length of fibres (LOF) and number of the fibres (NOF) connecting each pair of ROIs. Considering that the number of fibres is partially dependent on the surface area of their connected ROIs, we derived a new measure, the fibre density (FD), for the following association analysis, which is the fraction between number of fibres and the average surface of grey-matter regions i and j. These fibre density values are then used as the edge weights of the connectome network for each subject. The connectome network is now modelled as an undirected graph. We calculate the degree of each node (ROI) as the sum of weights of all connected edges to this node. This is the initial feature set for our deep learning analysis experiment.

In our study, we use the brain connectome network data from 163 subjects in 3 categories: HC (Healthy Control, 43 subjects); MCI (Mild Cognitive Impaired, 42 subjects) and AD (Alzheimer's Disease, 19 subjects). The brain connectome network data from the 163 subjects are 163 graphs. Since all the subjects share the same parcellation, they have the same set of node labels (the brain ROIs) but different connectivity (edges) between the nodes as the fibre densities are different for different subjects between the same pairs of ROI labels.

IV. THE DEEP LEARNING VISUALIZATION FRAMEWORK

The brain data classifier will first be trained using a training set that contains brain connectome networks of different subjects in 3 categories. We use a simple feature set of node degrees in the connectome networks as the input features of the deep neural network. Therefore, each input node represents a brain region of interest (ROI) and its network degree (connectivity). The visualization system can be activated when applying the trained model to a test subject which will generate activation values for all nodes in the hidden layers.

A. The Visualization Framework

The overall design of our deep learning visualization system is shown in Figure 1. The system has three visual components:

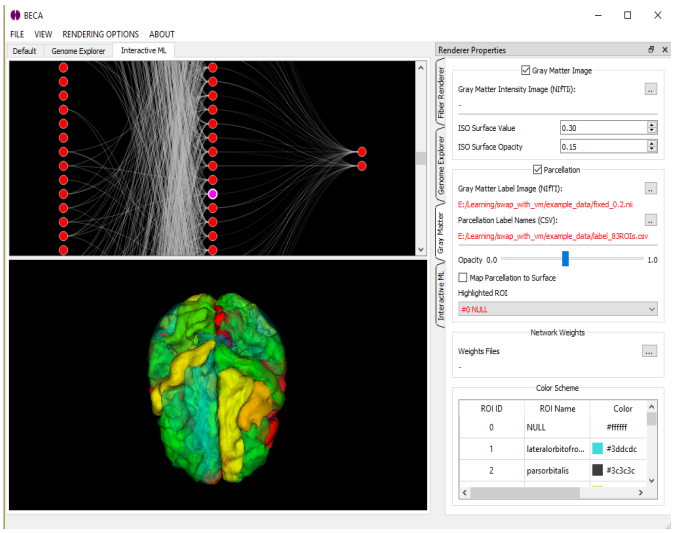

Figure 1. The interactive visualization system interface

(1) Network structure window. It shows the structures of the different layers of the deep neural network. All layers are drawn with their nodes and the edges connecting nodes between neighboring layers. The intensity / transparency of these edges are weighted by their edge weights. Zooming and scaling are implemented to allow the user to explore large networks. This window provides a structural overview of the learning process. It allows user to pick any of the nodes in the hidden layers to show the corresponding brain regions and their weights, calculated by our *importance* backpropagation algorithm. The user can also interactively edit and manipulate nodes and edges to make changes to the model. The interaction part will be further discussed in Section 4.3.

(2) Rendering window. This is where the 3D anatomical structure of a brain will be rendered. The regions of interest are segmented, and can be rendered differently. For a given (picked) node, the weights of the different brain regions associated with this node will be color-coded in the rendering window. 3D Rotation and zooming are implemented to allow 3D interactions with the brain rendering.

(3) Control window. This is a sub-window that controls various parameters for the visualization system such as intensity and transparency mapping, color mapping and thresholding.

Figure 2 shows the overall system architecture and flow. This chart only shows the interactive visual exploration part of the deep learning process. The formal cross-validation process is not shown in this flowchart, as cross-validation is carried out after the visual exploration stage. Thus, we can consider these visual test subjects an internal test set that is used to improve the performance of the learning model before cross-validation is applied.

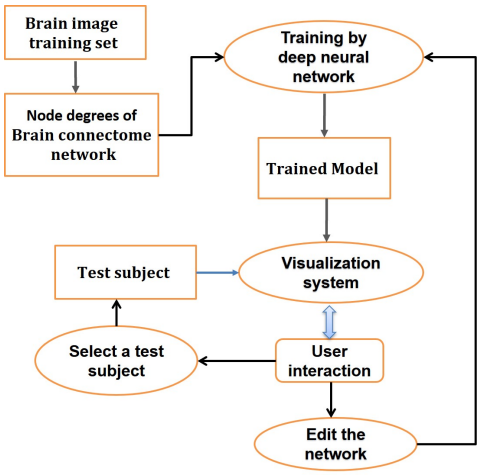

Figure 2. Overall system flowchart

B. Importance Factor Backpropagation

Once a deep learning model is trained, the user can run the model on a given test subject as input to observe how the model works with a certain type of input subject. Within our visualization system, the user can pick any node in the intermediate layers to show what the node represents in terms of weights of different brain regions. This provides the required contextual information to interpret the intermediate hidden layers. For example, when a brain region is consistently weighted high by some nodes in every layer, this could be an indication that this region is an important image phenotype for the Alzheimer's Disease classification problem. On the other hand, when a certain brain region is known to be associated with Alzheimer's Disease (e.g. the Hippocampus), nodes that carry high weights on this region may have more significance in the hidden layers, and could be a candidate to manipulate (e.g. increasing associated edge weights) when there is a need to re-train the mode to improve performance.

To compute the weights of brain regions for some node *S* (we will call it source node) in a hidden layer for a given input subject, we use a recursive method to backpropagate the *importance* factors from any node in a hidden layer back to the nodes in the input layer, which represents the network degrees of the different regions of interest (ROIs).

Let N_i^k denotes the *i*'s node in the *k*'s layer, with layer number 0 being the input layer. Let W_{ij}^k be the edge weight between node N_i^k and node N_j^{k-1}, and A_i^k be the activation value of node N_i^k (for the given input subject). We define the *importance* factor of node N_j^{k-1} (relative to the source nde) as:

$$r_j^{k-1} = \sum_i R_{ij}^{k-1}$$

where R_{ij}^{k-1} is the *importance* factor along edge N_j^{k-1} N_i^k, and is calculated by:

$$R_{ij}^{k-1} = \frac{|W_{ij}^k| * A_j^{k-1}}{\sum_l (|W_{il}^k| * A_{jl}^{k-1})} \times r_i^k$$

For example, the *importance* factor of node I1 along the edge J1I1 in Figure 3 is calculated as:

$$\frac{|w_1| * a_1}{|w_1| * a_1 + |w_2| * a_2 + |w_3| * a_2} \times r_{j1}$$

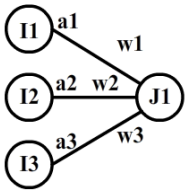

Figure 3. An example of the *importance* backpropagation.

Intuitively, the *importance* factor of a node represents how much this node contributes to the activation of the source node *S*. Starting from a source node *S*, the algorithm recursively computes an *importance* factor for every node in all possible paths from S back to the input layer. When a node can be reached by multiple paths, multiple *importance* factors will be computed, and their sum will be the final *importance* factor for this node. This process is recursively carried out backwards until reaching the input layer. At that point, we will have all r_i^0, which are the weights of all input brain regions for the given source node *S*. It is easy to prove that the sum of all final *importance* factors of all nodes in any given layer is 1.

C. Interactions

The interactive component of our visualization framework consists of two types of interactions: (1) brain structure exploration; and (2) editing of the deep neural network structure.

In brain structure exploration, the user will pick a node in a hidden layer. The brain rendering window will show the color-coded 3D rendering of the brain. Rotation and zooming can be used for the user to examine and understand which brain regions are activated by this node. By examining multiple nodes in different layers, the user will gain understanding about how different brain regions are affected by the evolution of the neural networks to draw a conclusion for different labelled test subjects (e.g. an AD patient vs a control). It may also provide clues about why certain test subjects are classified incorrectly. This could potentially lead to changes to the network structure for re-training.

The edit functions are designed to allow the user to make changes to nodes and edges based on user observations and domain knowledge. After changes are made, the deep neural network model will need to be re-trained using the revised network structure and weights as the new initial condition. Three types of edit operators are defined in our system:

- Deleting a node
- Adding a new node
- Changing the weight of an edge

These operations can be carried out interactively through combinations of user mouse clicking and control panel editing (Figure 4a). As an example, when a particular node has strong weight on a brain region that is known to be associated with Alzheimer's Disease, the user may decide to increase the weights of the edges that lead this node to the next layer.

Another scenario is when the model does not perform well with internal test set, we may increase the number of nodes or even layers to generate more feature combinations.

From an optimization point of view, these changes can be considered an educated perturbation of the initial conditions of learning process. When a modified neural network structure is used as the initial condition, it can potentially move the system out of a local minimum.

D. 3D Brain Image Rendering

For brain image rendering, we applied a traditional surface rendering technique on the brain regions of interest (ROIs). These surfaces are generated through brain image parcellation and iso-surface extraction (Figure 4b). As this is a 3D rendering, the user may rotate and zoom to gain the spatial perception of the associated regions for a picked node. For some internal regions that are hidden by other regions, it is also necessary to change the transparency of the surfaces to see through the outside layers to identify the interior regions.

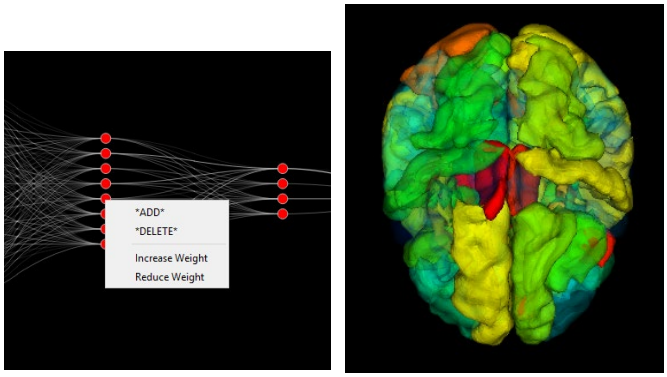

(a) (b)

Figure 4. (a) Node-editing interface; (b) Brain rendering with ROI labels

V. EXPERINEBTAL RESULTS

As a pilot study, we trained a three-layer fully connected neural network as a classifier to distinguish healthy controls from patients diagnosed with Alzheimer's Disease (AD). These fiber density values are then used as the edge weights of the connectome network for each subject. The connectome network is now modelled as an undirected graph. We calculate the degree of each node (ROI) as the sum of weights of all connected edges to this node. This is our initial feature set for our deep learning analysis experiment.

In our study, we focus on the analysis of the brain connectome data from 62 subjects in 2 categories: HC (Healthy Control, 43 subjects) and AD (Alzheimer's Disease, 19 subjects). The brain connectome network data from the 62 subjects are 62 graphs. Since all the subjects share the same parcellation, they have the same set of node labels (the brain ROIs) but different connectivity (edges) between the nodes as the fiber densities are different for different subjects between the same pairs of ROI labels.

The inputs consist of 62 subjects. Each subject is a vector of 83 elements representing the fiber densities of 83 regions of interest (ROIs) in the brain. We built an interactive exploration system that allows user to click any of the nodes in the neural network to show the corresponding features with respect to the input layer calculated by our proposed *importance* backpropagation algorithm. We also mapped the *importance* factor of each input ROI on a 3D brain, to help neuroscientist to have an intuitive understanding of which ROIs are correlative with each class. The discriminative ROIs for class AD are shown in Figure 4b. By visualizing the calculated *importance* of each ROI with a heatmap, we can see there are a small subset of the ROIs that are correlative with AD class, which are highlighted in red (Figure 5).

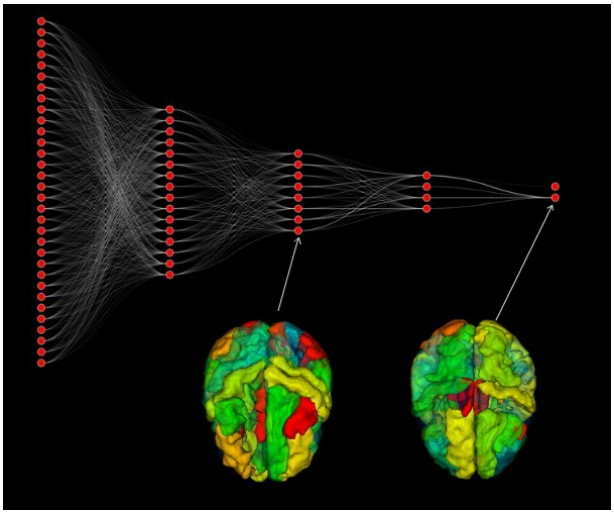

Figure 5. Brain ROIs (red) that are correlated to the AD class.

We also demonstrate the ability to change the neural network structure by interactively adding or deleting nodes with the system. As shown in Figure 6(a), we start with a three-layer fully connected neural network with 4 nodes in the second layer. After adding 4 extra nodes to the 2nd layer, as shown in Figure 6(b), the classification accuracy improves from 62.2% to 64.7% after retraining. We then add an extra layer of 4 nodes right before the output, as shown in Figure 6(c), which further improves the classification accuracy to 67.8%.

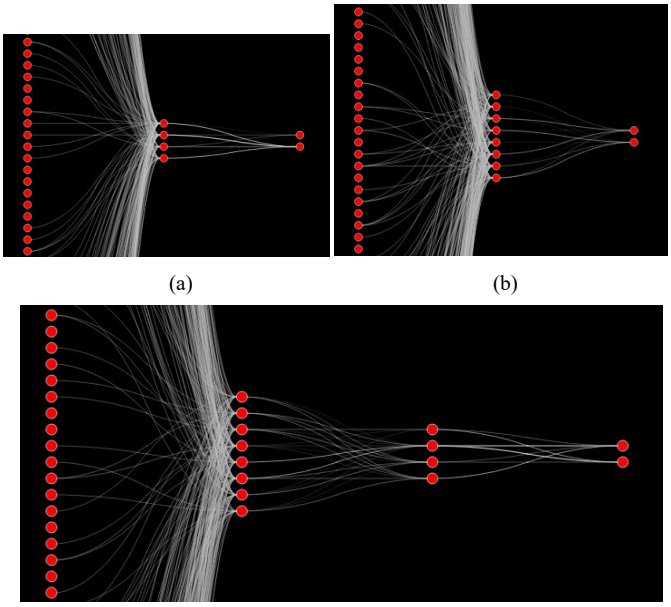

(a) (b)

(c)

Figure 5. Interactive editing of nodes and layers.

VI. CONCLUSIONS

We have presented an interactive visualization framework for a deep neural network system with brain image data. We have tested the system on a brain classification problem for the diagnosis of Alzheimer's Disease using a real world ADNI dataset. While the concept of visualizing deep neural networks to help user understand the learning model has been discussed before, we believe the key is to provide contextual information specific for the underlying applications. Our work focuses on building visual associations between deep neural network nodes and the complex brain structures. The backpropagation algorithm for computing the *importance* factors is efficient and easy to implement. Our system provides the needed biological and anatomical interpretation of the evolution of the deep neural network. Furthermore, it allows the users to make changes to the hidden layer nodes and edges based on their visual observations and domain knowledge to potentially improve the system's performance. This can be particularly useful for applications that require highly specialized knowledge and domain expertise (e.g. biomedical applications). User interference and re-direction can also be helpful when automatic deep learning performs poorly.

Although this paper focuses on the brain data classification problem, the visual analytics framework proposed here can be applied to a wide range of potential applications using deep learning techniques. In the future, we would like to develop more sophisticated interaction tools to facilitate better knowledge input and user feedback. We would also like to work with domain experts to collect more user experience to better evaluate the benefits of such a system.

ACKNOWLEDGMENTS

This work was supported in part by the NIH grant R01-EB022574.

REFERENCES

[1] Zachary C. Lipton. The mythos of model interpretability. ACM Queue, 16(3):31–57, June 2018.

[2] Thomas, J., Cook, K.: Illuminating the Path: Research and Development Agenda for Visual Analytics. IEEE-Press (2005)

[3] Amershi, S., Cakmak, M., Knox, W.B. and Kulesza, T., 2014. Power to the people: The role of humans in interactive machine learning. AI Magazine, 35(4), pp.105-120.

[4] Zahavy, T., Ben-Zrihem, N., Mannor, S. 2016. Graying the black box: Understanding dqns. In: ICML pp. 1899–1908.

[5] Rauber, P.E., Fadel, S., Falcao, A., Telea, A., 2017. Visualizing the hidden activity of artificial neural networks. IEEE TVCG 23 (1), 101–110.

[6] Tzeng, F.Y., Ma, K.L. 2005. Opening the black box - data driven visualization of neural networks. In: IEEE Visualization, pp. 383–390. http://dx.doi.org/10.1109/VISUAL.2005.1532820.

[7] Harley, A.W., 2015. An interactive node-link visualization of convolutional neural networks. In: International Symposium on Visual Computing. Springer, pp. 867–877.

[8] Streeter, M.J., Ward, M.O., Alvarez, S.A., 2001. Nvis: An interactive visualization tool for neural networks.

[9] Liu, M., Shi, J., Li, Z., Li, C., Zhu, J.J.H., Liu, S., 2017. Towards better analysis of deep convolutional neural networks. IEEE TVCG 23 (1), 91–100. http://dx.doi.org/10.

[10] Jason Yosinski, Jeff Clune, Anh Nguyen, Thomas Fuchs, and Hod Lipson. Understanding Neural Networks Through Deep Visualization. ICML Workshop on Deep Learning, 2015.

[11] Luisa M Zintgraf, Taco S Cohen, Tameem Adel, Max Welling. Visualizing Deep Neural Network Decisions: Prediction Difference Analysis. International Conference on Learning Representations (ICLR) 2017.

[12] Junpeng Wang, Liang Gou, Han-Wei Shen, and Hao Yang. DQNViz: A Visual Analytics Approach to Understand Deep Q-Networks. IEEE Trans Vis Comput Graph. September, 2018.

[13] Minsuk Kahng , Nikhil Thorat, Duen Horng Chau, Fernanda Viégas, Martin Wattenberg. GAN Lab: Understanding Complex Deep Generative Models using Interactive Visual Experimentation, IEEE Transactions on Visualization and Computer Graphics, 25(1) 2019, pp 310 – 320.

[14] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learningdeep features for discriminative localization. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2921–2929, June 2016.

[15] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Ba-tra. Grad-cam: Visual explanations from deep networks via gradient-basedlocalization. In 2017 IEEE International Conference on Computer Vision(ICCV), pages 618–626, Oct 2017.

[16] Ren, D., Amershi, S., Lee, B., Suh, J., Williams, J.D., 2017. Squares: Supporting interactive performance analysis for multiclass classifiers. IEEE TVCG 23 (1), 61–70.

[17] Alsallakh, B., Hanbury, A., Hauser, H., Miksch, S., Rauber, A., 2014. Visual methods for analyzing probabilistic classification data. IEEE TVCG 20 (12), 1703–1712.

[18] Chuang, J., Gupta, S., Manning, C.D., Heer, J. 2013. Topic model diagnostics: Assessing domain relevance via topical alignment. In: ICML, pp. 612–620.

[19] Shixia Liu, Xiting Wang, Mengchen Liu, Jun Zhu. Towards better analysis of machine learning models: A visual analytics perspective. Visual Informatics 1 (2017) 48–56.

[20] J Krause, A Perer, E Bertini. Using Visual Analytics to Interpret Predictive Machine Learning Models. ICML Workshop on Human Interpretability in Machine Learning, 2016.

[21] Krause, Josua, Perer, Adam, and Ng, Kenney. Interacting with predictions: Visual inspection of black-box machine learning models. ACM CHI 2016, 2016.

[22] Wang, J; Fang, S; Li, H; Goni, J; Saykin, AJ; Shen, L. Multigraph Visualization for Feature Classification of Brain Network Data. EuroVis Workshop on Visual Analytics (EuroVA), pp.61-65, 2016.

[23] Paiva, J.G.S., Schwartz, W.R., Pedrini, H., Minghim, R., 2015. An approach to supporting incremental visual data classification. IEEE TVCG 21 (1), 4–17.

[24] Liu, M., Liu, S., Zhu, X., Liao, Q., Wei, F., Pan, S., 2016. An uncertainty-aware approach for exploratory microblog retrieval. IEEE TVCG 22 (1), 250–259.

[25] Choo, J., Lee, C., Reddy, C.K., Park, H., 2013. Utopian: User-driven topic modeling based on interactive nonnegative matrix factorization. IEEE TVCG 19 (12) 1992–2001

[26] Wang, X., Liu, S., Liu, J., Chen, J., Zhu, J., Guo, B., 2016. TopicPanorama: A full picture of relevant topics. IEEE TVCG 22 (12), 2508–2521.

[27] Jiawei Zhang, Yang Wang, Piero Molino, Lezhi Li and David S. Ebert. Manifold: A Model-Agnostic Framework for Interpretation and Diagnosis of Machine Learning Models. IEEE Transactions on Visualization and Computer Graphics, 25(1), 2019, pp 364 – 373.

[28] Dominik Sacha, Matthias Kraus, Daniel A. Keim, and Min Chen . VIS4ML: An Ontology for Visual Analytics Assisted Machine Learning. IEEE Transactions on Visualization and Computer Graphics, 25(1), pp. 385 – 395.

[29] Li H, Fang S, Mukhopadhyay S, Saykin A, Shen L. Interactive machine learning by visualization: A small data solution. HMData'18: The 2nd IEEE Workshop on Human-in-the-loop Methods and Human Machine Collaboration in BigData. Seattle, WA, December 10, 2018。

[30] Manjón, J.V., Coupé, P., Concha, L., Buades, A., Collins, D.L. and Robles, M. Diffusion weighted image denoising using overcomplete local pca. PloS one, Vol. 8, No. 9, 2013, pp.e73021

[31] Cook, P., Bai, Y., Nedjati-Gilani, S., Seunarine, K., Hall, M., Parker, G. and Alexander, D. Camino: open-source diffusion-mri reconstruction and processing. 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine, Vol. 2759, Seattle WA, USA, 2006.

[32] Jenkinson M1, Beckmann CF, Behrens TE, Woolrich MW, Smith SM. FSL. Neuroimaging, 2012 Aug 15;62(2):782-90. doi: 10.1016.

[33] Shen, X., Tokoglu, F., Papademetris, X. and Constable, R.T. Groupwise whole-brain parcellation from resting-state fmri data for network node identification, Neuroimage. Vol. 82, 2013. pp.403–415.