

# NeuroMapper: In-browser Visualizer for Neural Network Training

Zhiyan Zhou\* Kevin Li\* Haekyu Park\* Megan Dass\* Austin Wright\* Nilaksh Das\* Duen Horng Chau \*



Figure 1: NEUROMAPPER visualizes how a deep neural network (DNN) model, Resnet-50 trained on CIFAR10 dataset in this example, evolves during training by displaying the embeddings for training samples at the model’s blocks throughout the training regiment. **A. Epoch Controller** for updating the embedding display by adjusting the training epoch through a slider or auto-incrementation using the controller. **B. Sample Size Controller** changes the number of samples. **C. Class Display Controller** toggles classes’ sample visibility. **D. AlignedUMAP Embedding View** visualizes input embeddings represented by each block in the DNN model, with user-adjustable hyperparameters at the bottom. **E. Example Embedding Evolution** of the fourth residual block of the model at epochs 10, 50, and 200.

## ABSTRACT

We present our ongoing work NEUROMAPPER, an in-browser visualization tool that helps machine learning (ML) developers interpret the evolution of a model during training, providing a new way to monitor the training process and visually discover reasons for suboptimal training. While most existing deep neural networks (DNNs) interpretation tools are designed for already-trained model, NEUROMAPPER scalably visualizes the evolution of the embeddings of a model’s blocks across training epochs, enabling real-time visualization of 40,000 embedded points. To promote the embedding visualizations’ spatial coherence across epochs, NEUROMAPPER adapts *AlignedUMAP*, a recent nonlinear dimensionality reduction technique to align the embeddings. With NEUROMAPPER, users can explore the training dynamics of a Resnet-50 model, and adjust the embedding visualizations’ parameters in real time. NEUROMAPPER is open-sourced at <https://github.com/poloclub/NeuroMapper> and runs in all modern web browsers. A demo of the tool in action is available at: <https://poloclub.github.io/NeuroMapper/>.

**Index Terms:** Human-centered computing—Visualization—Visualization systems and tools;

## 1 INTRODUCTION

Deep neural networks (DNNs) are now widely used in many domains, increasing the needs for understanding how they make de-

cisions [7]. While some existing interpretation approaches use embeddings to help people interpret DNNs, they are often designed for providing insight within a single epoch [8].

To scalably encode relationships of features and their dynamic changes during training while keeping the representation of samples across adjacent training epochs similar, we adapt a recent dimensionality reduction algorithm called *AlignedUMAP* [1]. *AlignedUMAP* not only distills the features of data samples into a low dimensional representation but also aligns the input embedding across training epochs so that the embedding of the same input data at adjacent epochs are more likely to be projected to similar vicinity, such that model embedding changes induced by training may be more easily discovered through the visualization. Naively applying UMAP [4] (or other nonlinear methods like t-SNE) for each epoch independently might lead to significant misalignment of embeddings, resulting in *data-visual discrepancies* where underlying data changes may not be easily inferred from changes in the visualization [2].

In our work, we contribute the following:

- **NEUROMAPPER, a cross-platform, browser-based tool that helps machine learning (ML) developers visualize the training process of deep neural networks (DNNs).** NEUROMAPPER provides a new way to monitor and interpret the model evolution during training, going beyond only visualizing performance metrics (e.g., accuracy, F1-score) or interpreting already-trained models [5].
- **Open-sourced visualization that scales to 40,000 embedding points in real time.** NEUROMAPPER takes the advantage of WebGL technology to drastically improve the visualization’s rendering speed for large-scale embedding data. NEUROMAPPER’s open-sourced code-base and demo is available at <https://github.com/poloclub/NeuroMapper>.

\*<sup>1</sup>Georgia Institute of Technology. {zzhou406|kevin.li|haekyu|mdass3|apwright|nilakshdas|polo}@gatech.edu

- **Aligned visualizations for high-dimensional model embedding evolution.** Empowered by AlignedUMAP, NEUROMAPPER visualizes the evolution of data representation by aligning the embedding of the same data points across training epochs. AlignedUMAP is a nonlinear approach and has potential to better represent high-dimensional data points [6], providing a complementary technique to recent linear approaches [2]. Users can adjust AlignedUMAP’s hyperparameters in real time.

## 2 DESIGNING NEUROMAPPER

### Aligning Input Data Representation Across Training Epochs.

Visualizing embeddings over training epochs presents unique challenges, because the naive approach of independently generating and visualizing each epoch’s embedding could lead to vastly different-looking visualizations at different epochs, where one epoch’s embedding could be non-linearly scaled, rotated (or even flipped) very differently than another epoch’s [2]. Conceptually, since the embedding of a later epoch evolves from an earlier epoch’s, the two epochs’ embeddings should share some spatial resemblance. To accomplish such spatial coherence, we adapt AlignedUMAP [1] to align the embeddings of input in adjacent training epochs. For each layer and epoch, NEUROMAPPER first computes the high-dimensional embedding of each input sample using the activation vector in the layer of the model at the training epoch. NEUROMAPPER then uses AlignedUMAP to align the 2D projection of such high-dimensional embeddings so that the embedding of the same input data at adjacent training epochs are projected on a similar location. Specifically, AlignedUMAP optimizes a dimensionality reducer for each epoch at the same time while adding a constraint on the distance between the embeddings of the same data at adjacent epochs [4]. **Open-sourced In-browser Scalable Visualization.** Developed using modern web technologies such as React, MobX, and ScatterGL, NEUROMAPPER’s interface allows users to interpret how 10,000 samples from the CIFAR-10 dataset are encoded in a Resnet-50 model and how the embeddings evolve during the training process. Users can access the tool with any modern web-browser.

**User Interface Design and Usage Scenario.** The main view of NEUROMAPPER shows how data representation evolves during training by visualizing the aligned 2D projections of the input data across training epochs. By visualizing the evolution of data representation, users can interpret how a model is trained (e.g., whether it is trained well or not). For instance, as seen in Fig. 2, NEUROMAPPER may help inform users whether a model is best-fit, under-fit, or over-fit by using the **Embedding View** (Fig. 1D) to visualize the aligned embedding of all ten thousands of data samples at the fourth block of a Resnet-50 model. Classes of the input data are differentiated by their color, and their display can be toggled in **Class Display** (Fig. 1C). To inspect the embedding, users can zoom and pan in the view. When users hover over a data point in a layer, the same data point in all different layers is highlighted, and its class is displayed next to the highlighted points. With the **Epoch Controller** (Fig. 1A) that looks like a typical video player, users can control which epoch to visualize. Users can either use the auto-incrementation functionality by clicking on the start/pause button or manually adjust the epoch by dragging the slider or pressing the arrow buttons. As users walk through the training epochs, they could notice the difference in separation between best-fit and over-fit models. For the best-fit model, there are clear separations between input embeddings of different classes, which is expected as block 4 is close to the output layer. For the over-fitted model, there are multiple further separations within a cluster of the data of the same class, meaning the model may overly differentiate the data of the same class. To further inspect different aspects of model evolution, users can adjust the **AlignedUMAP Hyperparameters** (Fig. 1D) such as  $N$ -neighbors and  $Min$  Distance.  $N$ -neighbors are the number of neighbors that AlignedUMAP look at when learning the input embeddings [3]. For example, if users set

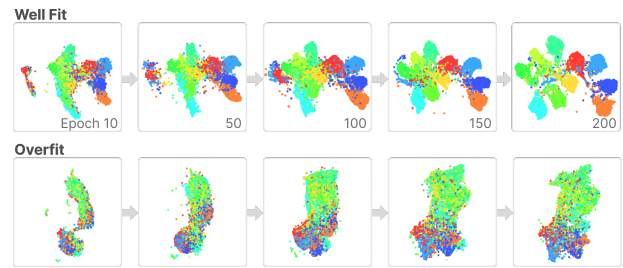


Figure 2: Evolution of embeddings at the fourth block of a well-fit and over-fit Resnet-50 model throughout training epochs.

lower values for  $N$ -neighbors, they can focus more on the changes in the fine details of the feature similarity between the input data instead of how the bigger picture of the data representation evolves; lower values of  $N$ -neighbors mean that AlignedUMAP focuses more on local neighborhoods of input data.  $Min$  Distance controls how tightly the embedding points are packed together [3]. Users can adjust the number of samples displayed in each panel through the **Sample Size Controller** (Fig. 1B)

## 3 CONCLUSION AND ONGOING WORK

NEUROMAPPER is an open-sourced interactive visualization tool for interpreting how intermediate representation of input evolve during training. Using WebGL, NEUROMAPPER can visualize large datasets in the modern browsers across multiple platforms. NEUROMAPPER leverages AlignedUMAP to show a continuous evolution of intermediate output during training. We are working on extending NEUROMAPPER to generate the input embedding in real time while a model is trained, so NEUROMAPPER runs alongside the training process. Our goal is to apply NEUROMAPPER to a wider range of models and datasets; it can handle visualizing tens of thousands of embedding points in real time, and we are planning to improve its scalability even more to handle more complex models and larger datasets through sampling. We also plan to improve NEUROMAPPER’s feature alignment by exploring Parametric UMAP, which utilizes a neural network to learn the relationship between data points and generate their embeddings.

## REFERENCES

- [1] Aligned umap basic usage. [https://umap-learn.readthedocs.io/en/latest/aligned\\_umap\\_basic\\_usage.html](https://umap-learn.readthedocs.io/en/latest/aligned_umap_basic_usage.html). Accessed: 2022-06-26.
- [2] M. Li, Z. Zhao, and C. Scheidegger. Visualizing neural networks with the grand tour. *Distill*, 2020. <https://distill.pub/2020/grand-tour>. doi: 10.23915/distill.00025
- [3] L. McInnes and J. Healy. Umap: Uniform manifold approximation and projection for dimension reduction. *ArXiv*, abs/1802.03426, 2018.
- [4] L. McInnes, J. Healy, N. Saul, and L. Grossberger. Umap: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 3(29):861, 2018.
- [5] D. Smilkov, N. Thorat, C. Nicholson, E. Reif, F. B. Viégas, and M. Wattenberg. Embedding projector: Interactive visualization and interpretation of embeddings. *arXiv preprint arXiv:1611.05469*, 2016.
- [6] L. van der Maaten and G. E. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008.
- [7] J. Wagner, C. M. Kohler, T. Gindele, L. Hetzel, J. T. Wiedemer, and S. Behnke. Interpretable and fine-grained visual explanations for convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9097–9107, 2019.
- [8] W. Zhong, C. Xie, Y. Zhong, Y. Wang, W. Xu, S. Cheng, and K. Mueller. Evolutionary visual analysis of deep neural networks. In *ICML Workshop on Visualization for Deep Learning*, p. 9, 2017.