

# *Pattern Recognition Letters*

## **Authorship Confirmation**

Please save a copy of this file, complete and upload as the “Confirmation of Authorship” file.

As corresponding author I, Suryaka Suresh, hereby confirm on behalf of all authors that:

1. This manuscript, or a large part of it, has not been published, was not, and is not being submitted to any other journal.
2. If presented at or submitted to or published at a conference(s), the conference(s) is (are) identified and substantial justification for re-publication is presented below. A copy of conference paper(s) is(are) uploaded with the manuscript.
3. If the manuscript appears as a preprint anywhere on the web, e.g. arXiv, etc., it is identified below. The preprint should include a statement that the paper is under consideration at Pattern Recognition Letters.
4. All text and graphics, except for those marked with sources, are original works of the authors, and all necessary permissions for publication were secured prior to submission of the manuscript.
5. All authors each made a significant contribution to the research reported and have read and approved the submitted manuscript.

Signature Suryaka Suresh

Date February 10, 2024

---

**List any pre-prints:** n/a

---

**Relevant Conference publication(s) (submitted, accepted, or published):** n/a

**Justification for re-publication:** n/a

### **Research Highlights (Required)**

- This paper categorizes the change in topological complexity through Betti-Numbers.
- Attempts to characterize the learning ability of the DNNs through topological transformation of feature space.
- Attempts to use the topological complexity of feature space for the task of ranking pre-trained models.



# On Characterizing the Evolution of Embedding Space of Neural Networks using Algebraic Topology

S. Suresh<sup>a,\*\*</sup>, B. Das<sup>b</sup>, V. Abrol<sup>a</sup>, S. Dutta Roy<sup>b</sup>

<sup>a</sup>Infosys Centre for AI & Department of Computer Science and Engineering, IIT Delhi, India

<sup>b</sup>Department of Electrical Engineering, IIT Delhi, India

Article history:

topology, deep learning, transfer learning

## ABSTRACT

We study how the topology of feature embedding space changes as it passes through the layers of a well-trained deep neural network (DNN) through Betti numbers. Motivated by existing studies using simplicial complexes on shallow fully connected networks (FCN), we present an extended analysis using Cubical homology instead, with a variety of popular deep architectures and real image datasets. We demonstrate that as depth increases, a topologically complicated dataset is transformed into a simple one, resulting in Betti numbers attaining their lowest possible value. The rate of decay in topological complexity (as a metric) helps quantify the impact of architectural choices on the generalization ability. Interestingly from a representation learning perspective, we highlight several invariances such as topological invariance of (1) an architecture on similar datasets; (2) embedding space of a dataset for architectures of variable depth; (3) embedding space to input resolution/size, and (4) data sub-sampling. In order to further demonstrate the link between expressivity & the generalization capability of a network, we consider the task of ranking pre-trained models for downstream classification task (transfer learning). Compared to existing approaches, the proposed metric has a better correlation to the actually achievable accuracy via fine-tuning the pre-trained model.

© 2024 Elsevier Ltd. All rights reserved.

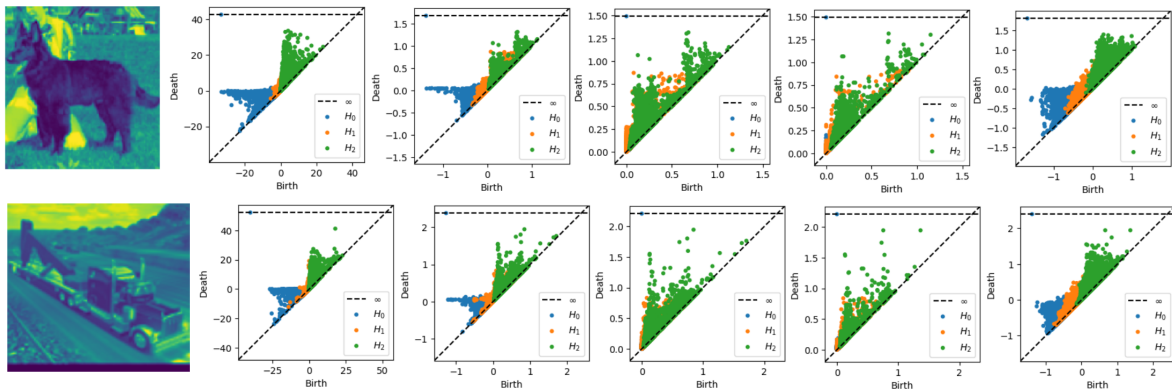
## 1. Introduction

Recently, deep learning has attracted a lot of attention in various learning tasks, as they offer great representational power as a result of non-linear transformation based layered architectures [1, 2]. A deep neural network's (DNN's) architectural components consist of affine-linear maps and nonlinearities, a combination of which is often termed as "black boxes" that lack explainability. While many properties such as expressivity, generalization error, sample/model complexity, and loss landscape of the

whole network have been studied extensively in various mathematical settings [3], the study of the evolution of feature embedding space has received little attention. Most existing studies linking expressivity and generalization in DNNs [3, 4, 5] emphasize the effect of depth on the learning dynamics in terms of Jacobian or Hessian of the network. However, such studies don't answer the question of which model is the most compatible for a dataset and task so as to provide practitioners valuable insights for empirical deep learning.

In this work, we aim to analyze the layer-by-layer embedding space of a DNN and leverage them to understand their learning characteristic empirically. Specifically, we use Algebraic topology [6] to capture high-order dependencies of embedding space in a data-dependent setting with popular pre-trained networks. Recently, topological tools

\*\*Corresponding author  
e-mail: [suryakas@iitd.ac.in](mailto:suryakas@iitd.ac.in) (S. Suresh),  
[bishshoy.das@ee.iitd.ac.in](mailto:bishshoy.das@ee.iitd.ac.in) (B. Das), [abrol@iitd.ac.in](mailto:abrol@iitd.ac.in)  
(V. Abrol), [sumantra@ee.iitd.ac.in](mailto:sumantra@ee.iitd.ac.in) (S. Dutta Roy)



**Fig. 1.** Persistence diagrams of embedding space of (from the left) Conv2d, BatchNorm2d, ReLU, MaxPool2d, and DenseBlock1, of DenseNet-169 for the images of the STL-10 dataset, respectively.  $H_i$  in the persistence diagram is the  $i^{\text{th}}$ -homology and homology is the algebraic object that defines the topology of a given space, here RGB image. Note how topological features at the block level are more interpretable compared to ones at individual layers like ReLU and MaxPool2d.

such as Persistent Homology [6], and K-mapper [7] have been widely explored to study DNNs, though restricted to shallow networks or low dimensional problems. We study the expressivity of a DNN in terms of extracted Betti numbers from the persistence diagram of the embedding space. Persistence homology maps complex data to persistence diagrams, a simple persistence-based feature vectorization that visualizes a concise summary of topological features in data. For statistical analysis, the persistence diagram needs to be equipped with a function that maps to this vector space like Betti numbers which is a quantification of the multiplicity of the homology. For a point cloud, Betti numbers represent a count of the topology features that appears in the subcomplex during filtration. The computational complexity involved in conventional algebraic topology tools is very high [4], and hence we present our analysis using Cubical homology [8] based on the method of cubical filtration that has been shown to be more efficient in extracting topological features of higher dimensional data like images. As an illustration, Fig. 1 shows the persistent diagrams corresponding to two example images from the STL-10 dataset of the embedding space of a DenseNet-169 model (pre-trained on ImageNet). One can attribute the similarities in the persistent diagrams to the network’s learning characteristics (feature extraction) from images while the dissimilarities to the image-specific attributes. Thus this provides a model-agnostic way of characterizing the embedding space via the quantification of persistent diagrams using Betti numbers. Experiments with a variety of popular deep architectures and real image datasets demonstrate that as depth increases, a topologically complicated dataset is transformed into a simple one, resulting in Betti numbers attaining their lowest possible value. Thus, we exploit the rate of decay in topological complexity  $\Omega$  [9, 10] (as a metric), an alternative to the popular Euler characteristic, to quantify the impact of architectural choices on the generalization ability. While doing so, we highlight and establish several interesting invariances such as the topological invariance of:

- an architecture on similar datasets.
- embedding of a dataset for an architecture of variable depth.
- embedding space to input resolution/size.
- embedding space to data sub-sampling.

We extend the work of Naitzat et al. (2020) [10] to multi-class higher dimensional problem settings. We leverage our analysis to establish a link between expressivity and the generalization capability of a network by considering the task of ranking pre-trained models for downstream supervised classification aka transfer learning. As explained earlier, we use the rate of decay of Betti numbers across layers as a metric via curve fitting and demonstrated an inversely proportional relationship between the slope and the actually achievable accuracy of a fine-tuned pre-trained network.

## 2. Related work

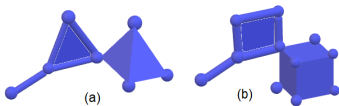
Carlsson et al. (2008) [11] and follow-up works from the group demonstrated how topological features serve as a potential measure of generalization in DNNs. However, their study was mainly restricted to visualization of the weight space of shallow networks (in terms of manifolds with well-known topology e.g., Klein bottle) trained on toy datasets such as MNIST using the Mapper algorithm [7]. Similarly, Rieck et al (2019) [12] used topology to track the evolution of DNN weights during training and proposed an early stopping criterion. In contrast, a series of works have attempted to study how DNNs transform a complex dataset to a topologically simpler one using Betti numbers. Bianchini et al. (2014) [13] proposed a theoretical measure to evaluate the complexity of DNN function space. Guss et al. (2018) [4] uses the Betti number as a measure of data complexity to characterize neural networks by the depth required to resolve the complexity. Hamada et al. (2018) [14] presented a data-driven test to establish if the transformed dataset is topologically simple. It doesn’t provide a quantification, rather just confirms if  $0^{\text{th}}$  order Betti number is greater than 1. Naitzat et al.

(2020) [10] explored persistent homology of the dataset as it passes through layers of a DNN. However, the study is restricted to shallow feed-forward networks and simulated dataset (binary labels) to prove the argument since the Betti number of a dataset are unknown in practice. Akai et al. (2021) [15] extended the work of [10] to study the stability of DNN mapping and the impact of outliers, but their analysis doesn't hold in high dimensional settings.

In contrast, work in this paper focuses on the representational power of a network with respect to a task (and dataset) at hand. We demonstrate the generalizability of a model through the evolution of the topological complexity of layer-by-layer embedding space. Compared to existing studies, the computational complexity around scaling a topological method to higher dimensional multi-class problems is tackled using recent advancements in Cubical homology. Finally, experiments in this work are done using a variety of popular architectures and real-world datasets.

### 3. Topological Preliminaries

In this section, we describe the relevant topological definitions based on the mathematical setting in [16].



**Fig. 2. Illustration of (a) simplicial complex and (b) cubical complex. The spheres represent the vertex i.e., a 0-dimensional simplex. Similarly, 1-simplex is an edge, 2-simplex is a triangle, and 3-simplex is a tetrahedron.**

#### 3.1. Persistence homology

Given a space  $X$  defined in terms of a point cloud  $\in \mathbb{R}^d$  with  $N$  points, the topological features can be computed using persistence homology. It requires the space to be represented using a simplicial complex i.e., constructing open sets around each data point using their intersection properties. This simplicial complex defines properties such as the number of connected components, holes, and voids it contains. Using a grid of scale values, a sequence of simplicial complexes is obtained to get information about the most persistent features that likely represent true underlying features of the shape of data. In order to work with image datasets (voxel data), recently cubical complexes have been shown to be efficient due to their robustness in higher dimensional problems [17]. These are equivalent to simplicial complexes with basic components such as squares, cubes, and higher dimensional analogs (see Fig. 2). The topological information is often summarized as a persistence diagram (see Fig. 1); a collection of points  $(x, y) \in \mathbb{R}^2$  representing a topological feature born at scale  $x$  and persisting until scale  $y$ . Points at infinity i.e. homology that never dies, are represented via an infinity bar in the diagram. Note:  $\beta_i$  is the  $H_i$  that exists at a threshold value.

#### 3.2. Quantification of Topology using Betti numbers

Using a persistent diagram, a topology can be quantified in terms of readily commutable Betti numbers  $\beta_k(X)$  that are the simplest invariants representing the shape of a manifold. For  $X \in \mathbb{R}^d$ , there are at most  $d$  non-zero Betti numbers. Typically, due to computational complexity constraints, the first three i.e.,  $0^{th}$  Betti number  $\beta_0(X)$ ,  $\beta_1(X)$  and  $\beta_2(X)$  that counts the connected components, holes, and voids are used. Readers are encouraged to refer to Section 2 from Naitzat et al. (2020) [10] for more details and examples. Here, the Betti curve is a function that maps the persistence diagram to an integer value over a range of threshold values  $\eta$ . Then  $\beta_k$  are computed for an appropriate value of  $\eta$ . We consider the embedding space of an image as the geometric object of interest, which is molded through homeomorphic or non-homeomorphic maps based on the task at hand. For instance, segmentation of an image would require a homeomorphic map in order to preserve the topology [18] while classification requires a non-homeomorphic map to collapse the space of images (each class) to a single point [10]. This work mainly focuses on the classification task and hence uses the Betti numbers, which is a unique signature for an object (here an image). Following [8] we use the voxel embedding space of an image  $x$  to compute  $\beta_k(x)$  which are combined to obtain a measure as:

$$\Omega = \frac{1}{N} \sum_{i=1}^N \omega_i \quad (1)$$

$$\omega_i = \beta_0(x_i) + \beta_1(x_i) + \dots + \beta_d(x_i); \quad x_i \in X$$

Here,  $\omega$  is the topological complexity of the embedding space of an image and  $\Omega$  is the expected topological complexity of the embedding space for a dataset. Note that eq(1) is a different measure than the one used in classical topology appearing in Morse theory [19] and adapted in [9, 10]. Later we empirically demonstrate that this measure is robust to data sampling as well as input resolution to a reasonable extent.

## 4. Experimental Setup

In this section, we provide details of the experimental protocol and software setup we used to carry out our experiments. We also delineate the different neural network architectures and datasets used in our study. All the experiments can be reproduced using the implementation available on GitHub<sup>1</sup>.

#### 4.1. Datasets

All the main experiments are done using a test set of three datasets (unless stated otherwise): STL-10 [20], CIFAR-10 [21] and Aircraft [22]. The STL-10 dataset

<sup>1</sup><https://github.com/Cross-Caps/DNNTopology>

consists of 500 training and 800 test images of  $96 \times 96$  resolution dispersed across 10 classes. The CIFAR-10 dataset consists of 50000 training and 10000 test images spread across 10 classes, with each image having a resolution of  $32 \times 32$ . Both STL-10 and CIFAR-10 pertain to the use cases where the target number of classes is small. Hence, we have also the test set of the Aircraft dataset consisting of 10200 images divided into 102 classes.

#### 4.2. Neural Network Architectures

1. VGG [23] is a sequential feedforward neural network with a uniform collection of convolutional layers followed by ReLU (Rectified Linear Unit) nonlinearity.
2. ResNet (or Residual Networks) [24] is a class of networks that features a skip connection that controls how much information flows to the next layer.
3. DenseNet [25] have transition layers that connect the dense blocks which spreads their weight function all over the neural network and allows feature reuse.
4. MobileNetV2 [26] is a lightweight neural network that uses depthwise separable convolutions. The residual blocks have a *wide*  $\rightarrow$  *narrow*  $\rightarrow$  *wide* approach.

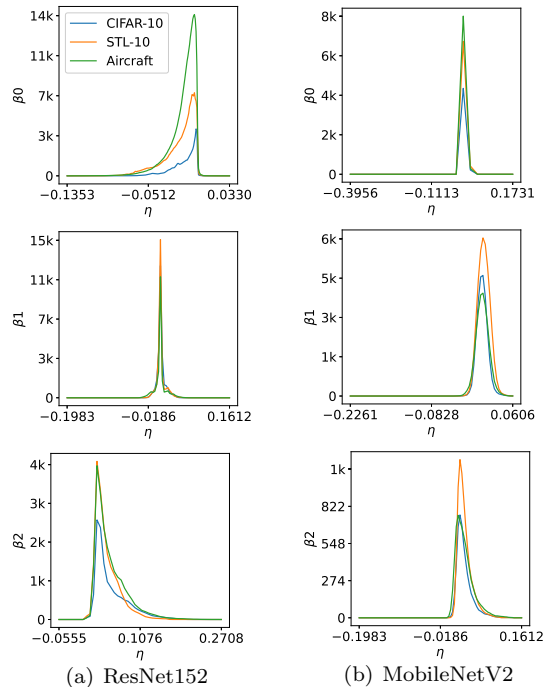
In total, we have experimented with 9 networks namely: VGG-16, ResNet-18, ResNet-50, ResNet-101, ResNet-152, MobileNetV2, DenseNet-121, DenseNet-169 and DenseNet-201. These models obtained from PyTorch model zoo are pre-trained on ImageNet dataset [27]. They use an input image size of  $256 \times 256$ , center cropped to  $224 \times 224$ , and coupled with standard data augmentations.

#### 4.3. Computation of Betti numbers

This work uses Cubical Ripser [8], built over the popular package Ripser [28], utilizing weighted cubical complexes to compute Betti numbers. The threshold value  $\eta$  is chosen such that the first three Betti numbers (at layer-1) are non-zero for all the images in the dataset. In the case of feedforward architectures like VGG, the measure  $\Omega$  is computed at each layer, while for complex architectures like ResNets/DenseNets, it is computed at each block after residual connection, considering the output volume data/tensor of shape  $(w \times h \times c)$  where  $w$  is the width,  $h$  is the height, and  $c$  is the number of channels.

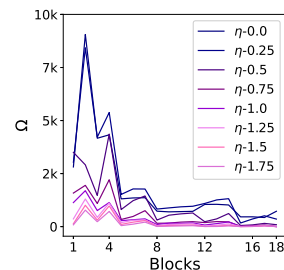
### 5. Topology of DNN Embedding Space using Cubical Complexes

As the inputs undergo layer-by-layer transformations in a DNN, a well-trained network should observe a decay in Betti numbers [10]. This is because with depth, a topologically complicated dataset is transformed into a simple one, resulting in Betti numbers attaining their lowest possible value. However, since as a DNN is expected to only focus and extract relevant information; a natural question arises: *if the topology induced by a pre-trained DNN is invariant to datasets?* To establish this for similar datasets consider Fig. 3 where one can clearly observe the near invariance of the



**Fig. 3. Betti curve for  $\beta_0$ ,  $\beta_1$  and  $\beta_2$  of the first layer embedding space of (a) ResNet-152 and (b) MobileNetV2 model, respectively.**

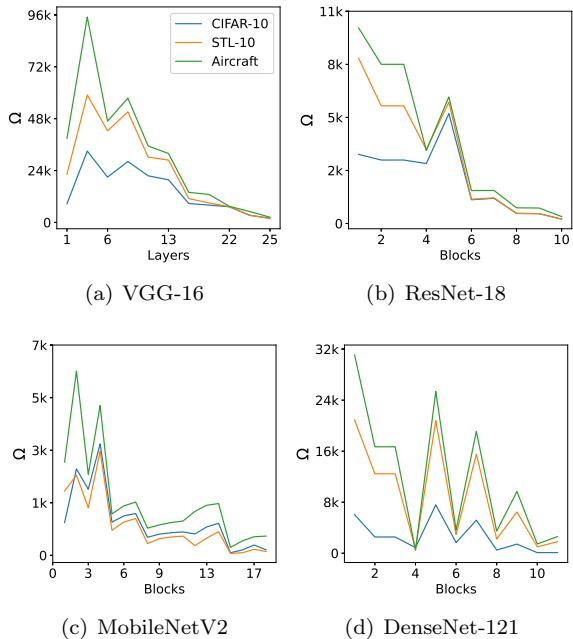
Betti curve (up to a scale) computed for the ResNet-152 model on different image datasets. This further illustrates the stability of the topological signature extracted by Betti numbers as a function of the chosen threshold.



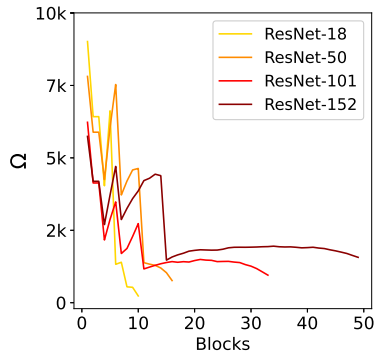
**Fig. 4. Topological complexity of embedding space in MobileNetV2 model for STL-10 dataset.**

Next, we explore the evolution of topological complexity of layer-by-layer embedding space via the proposed measure  $\Omega$ . Fig. 4 shows the measure  $\Omega$  as a function of the number of layers/blocks of a network. As expected due to non-homeomorphic maps (functions that don't preserve topology), we observe a decay in the value of  $\Omega$  with depth. One can categorize the layers/blocks into three distinct groups of blocks: 1) the initial ones which act as low-level feature extractors, 2) the middle ones with high-level features, and 3) the penultimate ones, which can be attributed to near identity maps where  $\Omega$  doesn't vary much. A notable observation is that  $\Omega$  is not sensitive to the change in the value of threshold  $\eta$  up to a scale that leads to robust interpretations about the learning





**Fig. 5.** Each figure show  $\Omega$  values for datasets **CIFAR-10**, **STL-10** and **Aircraft**. Each represents a different pattern and has the Betti numbers reduced across the layers.

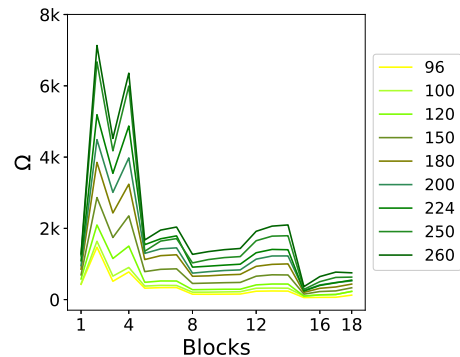


**Fig. 6.** Topological complexity of embedding space in ResNet model of variable depth for STL-10 dataset.

characteristics of a DNN. One can also observe the impact of architectural choices; for instance, the number of channels directly impacts the value of  $\Omega$ , and hence the decay in  $\Omega$  is not monotonic. As illustrated in Fig. 5 this behavior is consistent across datasets and network architectures. In particular, for the DenseNet model, we have also considered the transient block (with a pooling layer) in between two dense blocks separately. Observe how transient block (a contraction map) impacts and drastically reduces the value of  $\Omega$ . Hence, one should be careful while analyzing and interpreting  $\Omega$  across layers, blocks, or layers within a block, as the decay in value might not always be monotonic.

### 5.1. Impact of depth on topological complexity of the embedding space of an architecture

In the previous section, we establish the characterization of Betti numbers over depth for different architectures.



**Fig. 7.** Topological complexity of embedding space in MobileNetV2 model for STL-10 dataset at different input resize value.

However, in many tasks, the same architecture with variable depth is often compared for performance gains. In Fig. 6 we demonstrate the behaviour of  $\Omega$  for a pre-trained ResNet model of different depths. It can be observed that for a deeper architecture the number of blocks responsible for learning discriminative feature approach a limit (e.g., approx 15 for ResNet15 vs 10 for ResNet50). Also, the number of blocks for which  $\Omega$  is near constant increases with depth. This limiting is interesting from a representation learning perspective, as it shows that the amount of depth required for a downstream task is limited by the modeling capability of architecture and not its capacity. This is in line with existing studies that show that while theoretically deeper architectures are more expressive but for a given dataset not all layers contribute equally [29]. We observed similar characteristics in other architectures too, which establishes this important invariance.

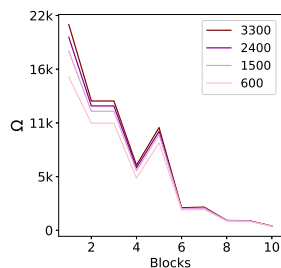
### 5.2. Impact of input resolution on the computation of topological complexity

In computer vision, pre-trained networks are often trained on inputs with a different resolution than the downstream task. For instance, transfer learning from the pre-trained network on the ImageNet dataset (with an image size of  $256 \times 256$ ) to the STL-10 dataset (with an image size of  $96 \times 96$ ). Hence, inputs are resized, such as by bilinear interpolation before feeding them to a pre-trained network. This makes working with topological tools on datasets with larger image sizes computationally expensive. We believe this might be the reason for earlier studies to restrict experimentation to either synthetic data or toy datasets like MNIST. Even though in the case of cubical homology, this cost increases exponentially with input size, it is near linear in the number of channels [8]. Due to the very nature of cubical filtration in computing homology, we argue that the behaviour of  $\Omega$  should be approx invariant to the input resolution. To demonstrate this, we computed the topological complexity of embedding space with inputs of variable size, and the results are reported in Fig. 7. It can be observed that  $\Omega$  is near invariant to input size up to a scale factor. This experiment suggests that for

pre-trained networks on ImageNet, input size in the range of 150 to 180 provides a good trade-off between computational complexity and explainability/interpretation of the learning dynamics of a network.

### 5.3. Impact of sample complexity on the computation of topological complexity

Similar to the input size, the categorization of embedding space using Betti numbers depends on the sample complexity [4]. In other words, more the data points better the estimate of topological complexity. However, the computational complexity can be prohibitive for large-scale datasets. Hence, in this experiment, we explore the impact of sample complexity on the computation of the topological complexity measure. For simplicity, we consider a balanced random subset of each class, although better sampling methods can also be explored. Results in Figure 8 demonstrate that the trajectory of  $\Omega$  with depth is near invariant to sample size. This has an important implication, as the experiments with different input & sample size suggest that a good trade-off can be achieved between computational complexity and quality of qualitative/quantitative results for DNN explainability using topology.



**Fig. 8. Topological complexity of embedding space in ResNet-18 model for Aircraft dataset with test set size ranging from 600 to 3300.**

**Table 1. Comparison of Pearson correlation coefficient ( $\rho$ ) of LEEP and the proposed TTP measure. Correlations are computed with respect to test accuracies in various experimental settings**

Experimental Setup			$\rho$	
Source	Target	Model	LEEP	TTP
ImageNet	STL-10	ResNet-18	.86	.94
		ResNet-50	.72	.70
		ResNet-152	.76	.87
		DenseNet-169	.68	.93
	Aircraft	ResNet-18	.92	.87
		ResNet-50	.82	.89
		DenseNet-169	.68	.81
	CIFAR-10	DenseNet-201	.78	.92
		ResNet-50	.91	.89
		ResNet-101	.73	.87
		MobileNetV2	.69	.64
	Caltech-101	DenseNet-169	.87	.95
ResNet-50		.67	.72	
		DenseNet-169	.58	.65

## 6. Case Study: Model Ranking for Transfer Learning

In this experiment, we establish a link between the topological complexity of a network and its generalization capability. To this aim, we consider the model ranking problem where given a target dataset and a large pool of pre-trained models, the goal is to rank the models in terms of the ease of knowledge transfer from the source model to the target dataset. This saves the computation cost involved with manually fine-tuning pre-trained models to select an optimal one. Existing studies in this space can be broadly classified into two categories with ones that compute: 1) a distribution over label space to find source and target label compatibility e.g., NCE [30], LEEP [31], Neural checkpoint [32]; and 2) a distribution over feature space to find last layer embedding and target label compatibility e.g., LogME [33], Model-Linearization [34]. We show that one can also use the trajectory or the rate of decay of the proposed measure  $\Omega$  to evaluate the transferability of representations learned by pre-trained classifiers. We denote this topological transferability measure (TTP) by  $\theta$ , computed as the instantaneous slope of the polynomial  $P$  fitted to  $\Omega$  values across the layers/blocks of the neural network i.e.,

$$\theta = \text{slope}(P)|_b, \quad (2)$$

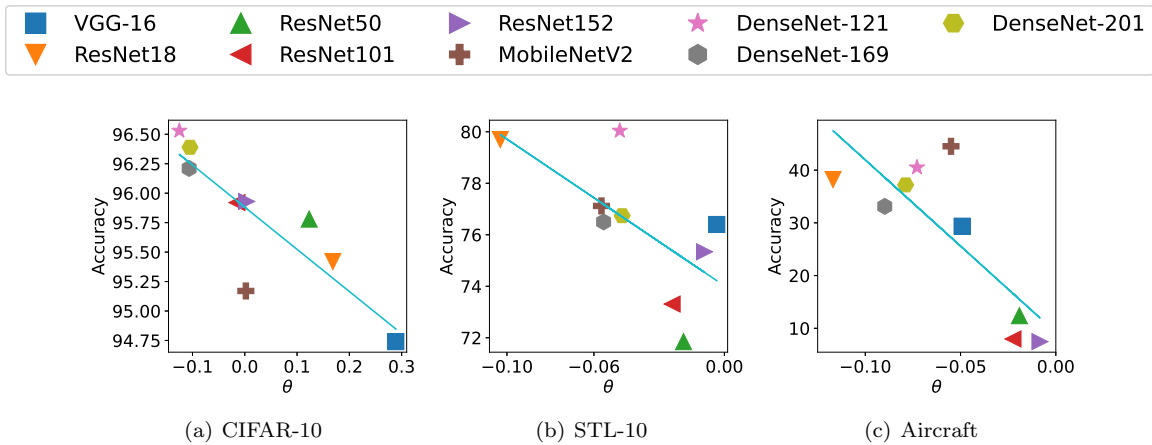
where,  $b$  is the midpoint on the curve  $P$ . The aim here is not to propose a state-of-the-art transferability measure but to demonstrate the power of topological tools for model interpretability. However, note that, unlike most existing measures,  $\theta$  can also be applied when the source and target tasks have different label semantics or are heterogeneous.

We perform transfer learning experiments with the datasets and models described in Section 4, where only the penultimate classification layers of the models were fine-tuned to establish the baselines. In Fig. 9 we show the negative correlation (inverse relationship) between the measure  $\theta$  and the maximum achievable accuracy if a network is actually fine-tuned on the dataset. Table 1 shows the correlation coefficients of the state-of-the-art LEEP measure in comparison with the proposed measure. It can be observed that the proposed measure performs comparably to LEEP for shallow to medium-depth models. However, the proposed measure consistently performs better in the case of very deep architectures. This is because the proposed measure considers the network as a whole, while LEEP only considers the last layer. Embedding from the penultimate layer has less discrimination ability than some intermediate layers due to the input size and complexity of the target dataset; hence, the initial layer with frozen weights is often used as a feature extractor in the transfer learning framework.

## 7. Conclusion

This work attempts to demonstrate the potential of topological tools in quantifying the complexity of embedding space in a DNN. In particular, we propose a measure





**Fig. 9.** Correlation between fine-tuned accuracy (Y-axis) and TTP measure ( $\theta$ ) (X-axis) for model ranking on 3 datasets with 9 pre-trained models.

based on Betti numbers to quantify the topological complexity. Experiments with a variety of architectures unravel interesting invariance which provides important cues for characterizing the learning behavior in DNNs. Further, we established a link between the expressivity & the generalization capability of a network by considering a case study on ranking pre-trained models for quantifying the ease of transfer learning from a source to a target dataset.

## 8. Acknowledgment

This work is supported by IIITD-IITD joint research grant (MFIRP-233/004) and Infosys Foundation via Infosys Centre for AI, IIIT Delhi.

## References

- [1] Y. Guo, H. Wang, Q. Hu, H. Liu, L. Liu, M. Bennamoun, Deep learning for 3d point clouds: A survey, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43 (12) (2021) 4338–4364. doi:10.1109/TPAMI.2020.3005434.
- [2] X.-S. Wei, Y.-Z. Song, O. M. Aodha, J. Wu, Y. Peng, J. Tang, J. Yang, S. Belongie, Fine-grained image analysis with deep learning: A survey, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44 (12) (2022) 8927–8948. doi:10.1109/TPAMI.2021.3126648.
- [3] B. Poole, S. Lahiri, M. Raghu, J. Sohl-Dickstein, S. Ganguli, Exponential expressivity in deep neural networks through transient chaos, in: *International Conference on Neural Information Processing Systems (NeurIPS)*, 2016, pp. 3368–3376. URL [https://proceedings.neurips.cc/paper\\_files/paper/2016/file/148510031349642de5ca0c544f31b2ef-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2016/file/148510031349642de5ca0c544f31b2ef-Paper.pdf),
- [4] W. H. Guss, R. Salakhutdinov, On characterizing the capacity of neural networks using algebraic topology, in: *NeurIPS workshop on Deep Learning: Bridging Theory and Practice*, 2017, p. 1=13. doi:10.48550/arXiv.1802.04443.
- [5] M. Murray, V. Abrol, J. Tanner, Activation function design for deep networks: linearity and effective initialisation, *Applied and Computational Harmonic Analysis* 59 (2022) 117–154, special Issue on Harmonic Analysis and Machine Learning. doi:10.1016/j.acha.2021.12.010.
- [6] F. Chazal, B. Michel, An introduction to topological data analysis: Fundamental and practical aspects for data scientists, *Frontiers in Artificial Intelligence* 4 (2021). doi:10.3389/frai.2021.667963.
- [7] G. Singh, F. Memoli, G. Carlsson, Topological Methods for the Analysis of High Dimensional Data Sets and 3D Object Recognition, in: M. Botsch, R. Pajarola, B. Chen, M. Zwicker (Eds.), *Eurographics Symposium on Point-Based Graphics*, The Eurographics Association, 2007, pp. 91–100. doi:10.2312/SPBG/SPBG07/091-100.
- [8] S. Kaji, T. Sudo, K. Ahara, Cubical ripser: Software for computing persistent homology of image and volume data (2020). doi:https://doi.org/10.48550/arXiv.2005.12692.
- [9] M. Bianchini, F. Scarselli, On the complexity of neural network classifiers: A comparison between shallow and deep architectures, *IEEE Transactions on Neural Networks and Learning Systems* 25 (8) (2014) 1553–1565. doi:10.1109/TNNLS.2013.2293637.
- [10] G. Naitzat, A. Zhitnikov, L.-H. Lim, Topology of deep neural networks, *Journal of Machine Learning Research* 21 (184) (2020) 1–40. URL <http://jmlr.org/papers/v21/20-345.html>
- [11] G. E. Carlsson, T. Ishkhanov, V. de Silva, A. Zomorodian, On the local behavior of spaces of natural images, *International Journal of Computer Vision* 76 (1) (2008) 1–12. doi:10.1007/s11263-007-0056-x.
- [12] B. Rieck, M. Togninalli, C. Bock, M. Moor, M. Horn, T. Gumbusch, K. Borgwardt, Neural persistence: A complexity measure for deep neural networks using algebraic topology, in: *International Conference on Learning Representations (ICLR)*, 2019, pp. 1–11. URL <https://openreview.net/forum?id=ByxkijC5FQ>
- [13] M. Bianchini, F. Scarselli, On the complexity of neural network classifiers: A comparison between shallow and deep architectures, *IEEE Transactions on Neural Networks and Learning Systems* 25 (8) (2014) 1553–1565. doi:10.1109/TNNLS.2013.2293637.
- [14] N. Hamada, K. Goto, Data-driven analysis of pareto set topology, in: *Genetic and Evolutionary Computation Conference*, 2018, p. 657–664. doi:10.1145/3205455.3205613.
- [15] N. Akai, T. Hirayama, H. Murase, Experimental stability analysis of neural networks in classification problems with confidence sets for persistence diagrams, *Neural Networks* 143 (2021) 42–51. doi:https://doi.org/10.1016/j.neunet.2021.05.007.
- [16] T. K. Dey, Y. Wang, *Computational Topology for Data Analysis*, Cambridge University Press, 2022. doi:10.1017/9781009099950.
- [17] S. Choe, S. Ramanna, Cubical homology-based machine learning: An application in image classification, *MDI Axioms* 11 (3) (Mar 2022). doi:10.3390/axioms11030112.
- [18] X. Hu, F. Li, D. Samaras, C. Chen, Topology-preserving deep

- image segmentation, in: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Vol. 32, Curran Associates, Inc., 2019, pp. 1–12.  
URL [https://proceedings.neurips.cc/paper\\_files/paper/2019/file/2d95666e2649f5ad9-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2019/file/2d95666e2649f5ad9-Paper.pdf)
- [19] J. W. Milnor, *Morse theory*, no. 51, Princeton university press, 1963.
- [20] A. Coates, A. Ng, H. Lee, [An analysis of single-layer networks in unsupervised feature learning](#), in: G. Gordon, D. Dunson, M. Dudík (Eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, Vol. 15 of *Proceedings of Machine Learning Research*, PMLR, Fort Lauderdale, FL, USA, 2011, pp. 215–223.  
URL <https://proceedings.mlr.press/v15/coates11a.html>
- [21] A. Krizhevsky, [Learning multiple layers of features from tiny images](#), Tech. rep., University of Toronto (2009).  
URL <https://www.cs.toronto.edu/~kriz/cifar.html>
- [22] S. Maji, J. Kannala, E. Rahtu, M. Blaschko, A. Vedaldi, [Fine-grained visual classification of aircraft](#), Tech. rep., University of Oxford (2013). [arXiv:1306.5151](#).  
URL <https://www.robots.ox.ac.uk/~vgg/data/fgvc-aircraft/>
- [23] K. Simonyan, A. Zisserman, [Very deep convolutional networks for large-scale image recognition](#), in: Y. Bengio, Y. LeCun (Eds.), *International Conference on Learning Representations, ICLR*, 2015, pp. 1–14. [doi:10.48550/arXiv.1409.1556](#).
- [24] K. He, X. Zhang, S. Ren, J. Sun, [Deep residual learning for image recognition](#), in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778. [doi:10.1109/CVPR.2016.90](#).
- [25] G. Huang, Z. Liu, L. Van Der Maaten, K. Q. Weinberger, [Densely connected convolutional networks](#), in: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2261–2269. [doi:10.1109/CVPR.2017.243](#).
- [26] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, L.-C. Chen, [Mobilenetv2: Inverted residuals and linear bottlenecks](#), in: *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4510–4520. [doi:10.1109/CVPR.2018.00474](#).
- [27] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, L. Fei-Fei, *ImageNet Large Scale Visual Recognition Challenge*, *International Journal of Computer Vision (IJCV)* 115 (3) (2015) 211–252. [doi:10.1007/s11263-015-0816-y](#).
- [28] U. Bauer, [Ripsper: efficient computation of vietoris-rips persistence barcodes](#), *Springer Journal of Applied and Computational Topology* (2021). [doi:10.1007/s41468-021-00071-5](#).
- [29] C. Zhang, S. Bengio, Y. Singer, [Are all layers created equal?](#), *Journal of Machine Learning Research* 23 (67) (2022) 1–28.  
URL <http://jmlr.org/papers/v23/20-069.html>
- [30] A. Tran, C. Nguyen, T. Hassner, [Transferability and hardness of supervised classification tasks](#), in: *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 1395–1405. [doi:10.1109/ICCV.2019.00148](#).
- [31] C. Nguyen, T. Hassner, M. Seeger, C. Archambeau, [LEEP: A new measure to evaluate transferability of learned representations](#), in: H. D. III, A. Singh (Eds.), *International Conference on Machine Learning (ICML)*, Vol. 119 of *Proceedings of Machine Learning Research*, PMLR, 2020, pp. 7294–7305.  
URL <https://proceedings.mlr.press/v119/nguyen20b.html>
- [32] Y. Li, X. Jia, R. Sang, Y. Zhu, B. Green, L. Wang, B. Gong, [Ranking neural checkpoints](#), in: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 2662–2672. [doi:10.1109/CVPR46437.2021.00269](#).
- [33] K. You, Y. Liu, J. Wang, M. Long, [Logme: Practical assessment of pre-trained models for transfer learning](#), in: *International Conference on Machine Learning (ICML)*, 2021, pp. 12133–12143.
- [34] A. Deshpande, A. Achille, A. Ravichandran, H. Li, L. Z. and-Charless Fowlkes, R. Bhotika, S. Soatto, P. Perona, [A linearized framework and a new benchmark for model selection for fine-tuning](#) (2021).  
URL <https://arxiv.org/abs/2102.00084>