FeaStNet: Feature-Steered Graph Convolutions for 3D Shape Analysis – Supplementary Material –

Nitika Verma Edmond Boyer Jakob Verbeek

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

firstname.lastname@inria.fr

1. Semi-supervised document classification

We experiment with semi-supervised document classification on the Cora and PubMed datasets [3]. These datasets contain scientific papers divided into seven and three classes respectively. The data is organized in a citation graph, where documents correspond to vertices, and citations are reflected by edges. Each document (2,708 in Cora and 19,717 in PubMed) is represented by sparse bag-of-words feature vector of dimension 1,433 for Cora and 500 for PubMed. There are 5,429 and 44,338 citation links in Cora and PubMed respectively. We followed the experimental setup and the dataset split as given in [1, 4]. There are 20 training samples per class and 500 and 1,000 vertices in the validation and test set respectively. We used two convolutional layers (Conv16 + Conv16). We used the validation data to choose the number of weight matrices M, learning rate, and the weight-decay level.

For M, we tried values from 1 upto 32, and found $M = 1$ to give best results on the validation data. This suggests that our model did not succeed in learning documents features to define filter shapes, and instead uses a uniform filter over the graph. The results in Table 1 show that the classification accuracy we obtain is comparable to the results of other recent graph convolutional approaches [1, 2], and significantly improves over the graph embedding approach of Yang *et al*. [4].

Method	Cora	PubMed
Planetoid [4]	75.7%	77.2%
GCN [1]	81.6%	78.7%
MoNet [2]	81.7%	78.8%
FeaStNet	81.6%	79.0%

Table 1. Classification accuracy on the Cora and PubMed document classification datasets.

2. Activation Visualization

The single-scale FeaStNet model is composed of a sequence of linear and graph convolution layers: Lin16+Conv32+Conv64+Conv128+Lin256+Lin6890 where the numbers indicate the amount of output features of each layer. In the figures given below, we show the activations of some randomly selected features learned across the layers given above.

Gradually across the layers the features become more pose invariant and more localized as required by the task. We use coordinates (xyz) as input for Figure 1 and Figure 2. The first linear layer here shows activations as linear functions of the coordinates which is not the case in Figure 3 and Figure 4 which use shot descriptors as input.

References

[1] T. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.

[2] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *CVPR*, 2017.

[3] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93, 2008.

[4] Z. Yang, W. Cohen, and R. Salakhutdinov. Revisiting semi-supervised learning with graph embeddings. In *ICML*, 2016.

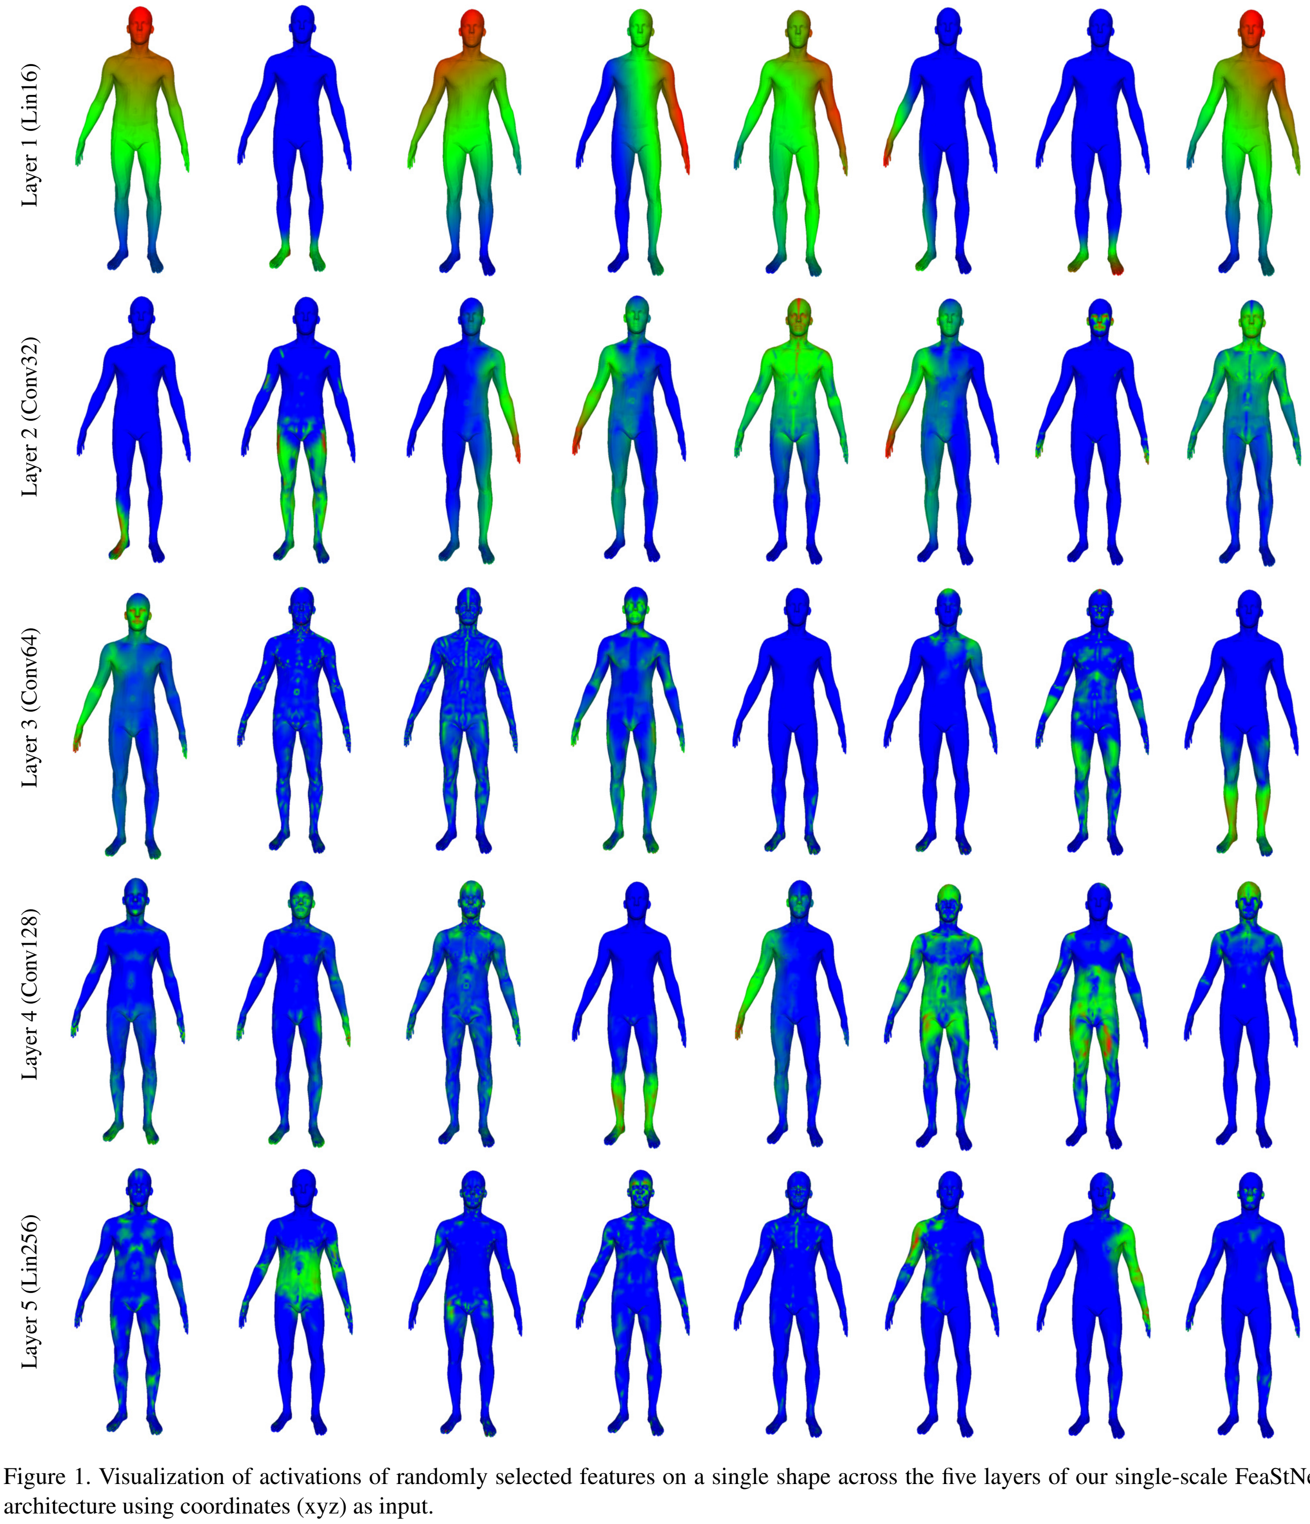

Figure 1. Visualization of activations of randomly selected features on a single shape across the five layers of our single-scale FeaStNet architecture using coordinates (xyz) as input.

Figure 2. Visualization of activations of randomly selected features on different shapes across the five layers of our single-scale FeaStNet architecture using coordinates (xyz) as input.

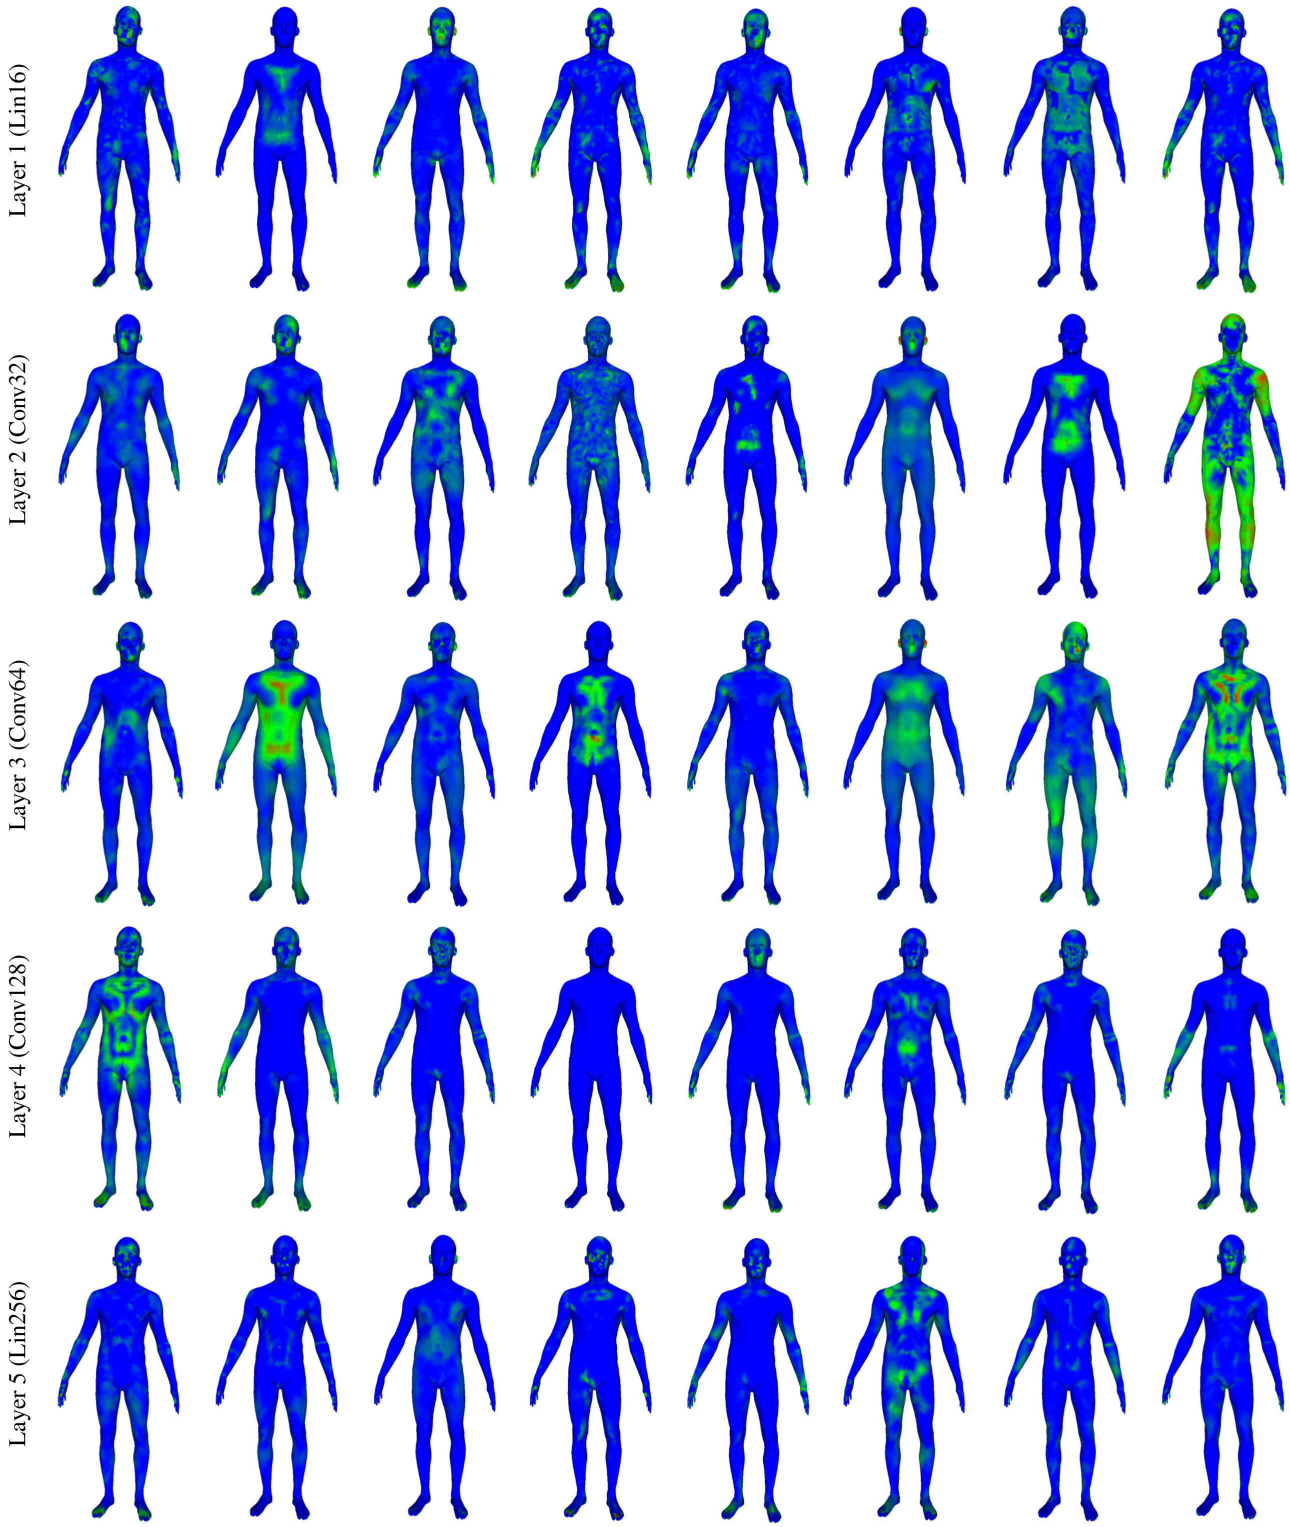

Figure 3. Visualization of activations of randomly selected features on a single shape across the five layers of our single-scale FeaStNet architecture using shot descriptors as input.

Figure 4. Visualization of activations of randomly selected features on different shapes across the five layers of our single-scale FeaStNet architecture using shot descriptors as input.