Graph-based Interpolation of Feature Vectors for Accurate Few-Shot Classification

Yuqing HU[1,2], Vincent GRIPON[1], Stéphane PATEUX[2]

[1]Lab-STICC, IMT Atlantique, Brest, France

[2]Orange Labs, Cesson-Sévigné, France

Abstract

In few-shot classification, the aim is to learn models able to discriminate classes using only a small number of labeled examples. In this context, works have proposed to introduce **Graph Neural Networks** (GNNs) aiming at exploiting the information contained in other samples treated concurrently, what is commonly referred to as the transductive setting in the literature. These GNNs are trained all together with a backbone **feature extractor**. In this paper, we propose a new method that relies on graphs only to interpolate feature vectors instead, resulting in a **transductive** learning setting with no additional parameters to train. Our proposed method thus exploits two levels of information : a) transfer features obtained on generic datasets, b) transductive information obtained from other samples to be classified. Using standard few-shot vision classification datasets, we demonstrate its ability to bring significant gains compared to other works.

Summary

- Given a task containing K_n classes and a small amount s of labelled images per class, the goal is to train a classifier that predicts the class of Q unlabelled inputs.
- To solve such problems, we introduce a three-stage method combining transfer learning, a graph-based interpolation technique and logistic regression.
- We empirically demonstrate that the proposed method reaches competitive accuracy on standardized benchmarks in the field of few-shot learning and largely surpasses the current works using GNNs.

Methodology

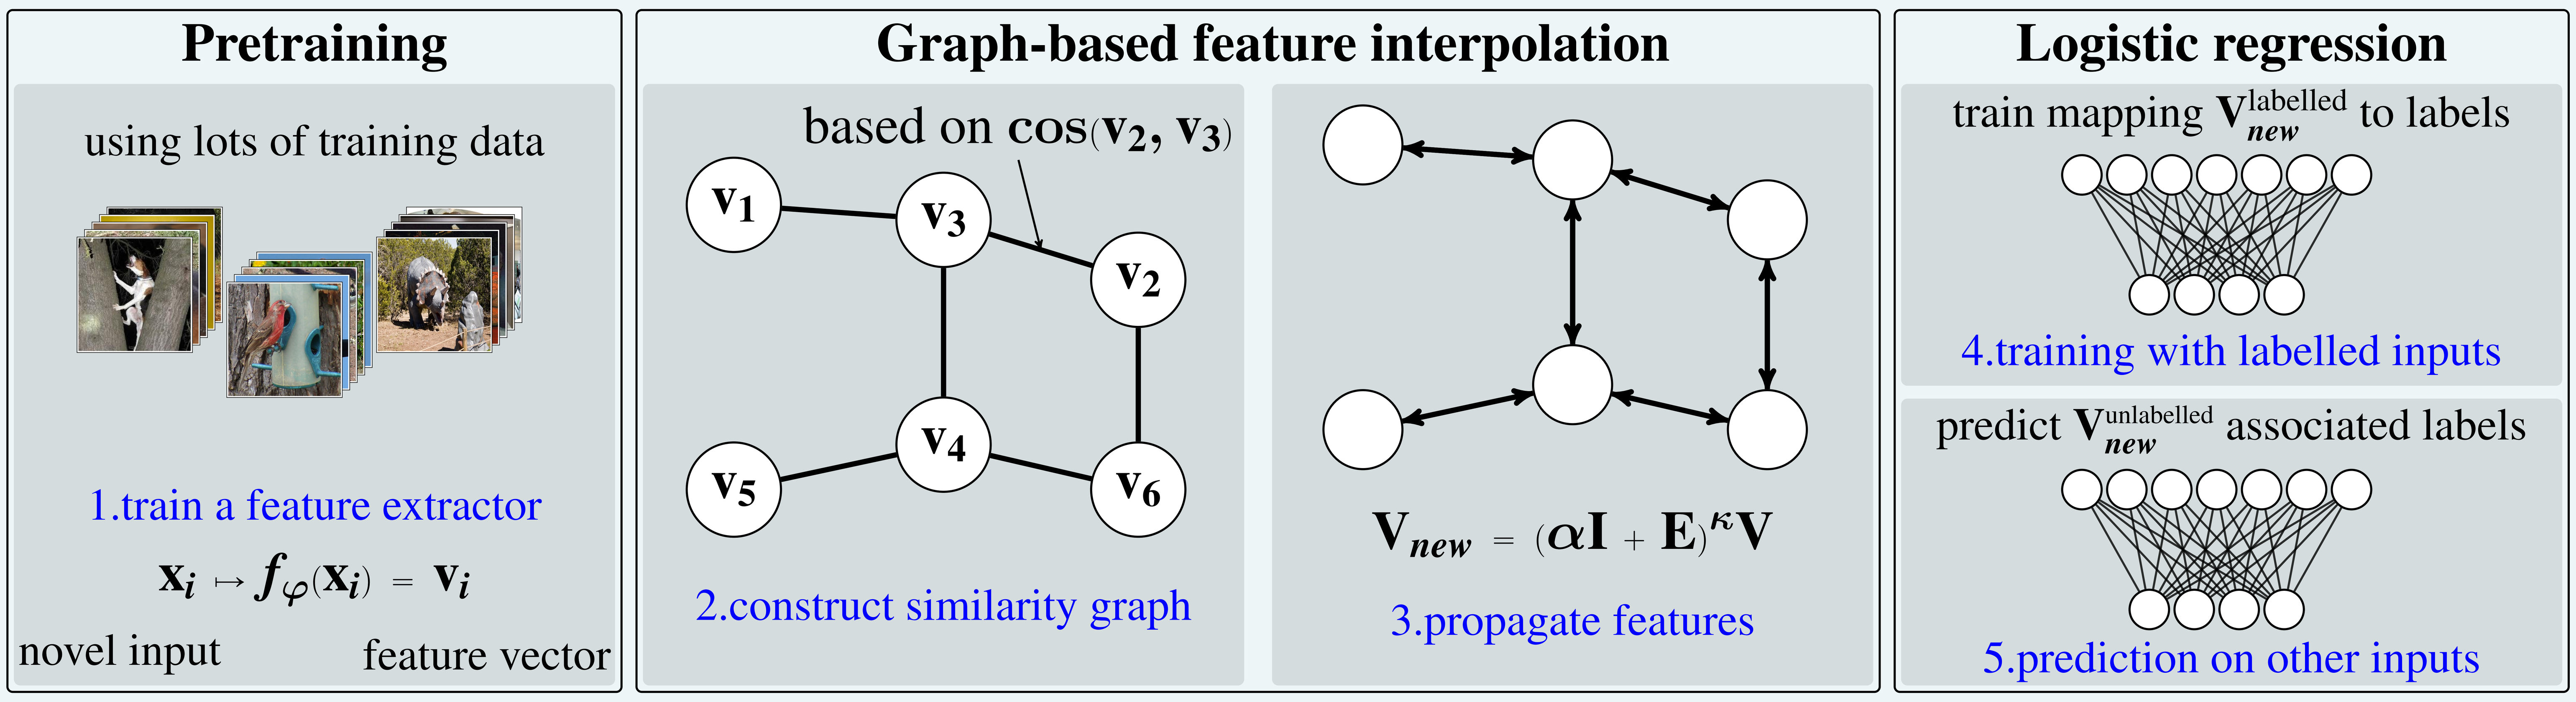

Results

Table 1 - 1-shot and 5-shot accuracy of state-of-the-art methods in the literature, compared with the proposed solution.

Method	Backbone	miniImageNet 1-shot	5-shot
MAML [1]	ResNet18	49.61 ± 0.92%	65.72 ± 0.77%
ProtoNet [5]	ResNet18	54.16 ± 0.82%	73.68 ± 0.65%
S2M2_R [4]	ResNet18	64.06 ± 0.18%	80.58 ± 0.12%
LaplacianShot [7]	ResNet18	72.11 ± 0.19%	82.31 ± 0.14%
Transfer+Graph Interpolation (ours)	ResNet18	**72.40 ± 0.24%**	**82.89 ± 0.14%**
ProtoNet [5]	WRN	62.60 ± 0.20%	79.97 ± 0.14%
S2M2_R [4]	WRN	64.93 ± 0.18%	83.18 ± 0.11%
SIB [2]	WRN	70.00 ± 0.60%	79.20 ± 0.40%
Transfer+Graph Interpolation (ours)	WRN	**76.50 ± 0.23%**	**85.23 ± 0.13%**

Table 2 - 1-shot and 5-shot accuracy on miniImageNet, when using various Graph Neural Networks.

Method	1-shot	5-shot
Transfer+GAT [6]	65.38 ± 0.89%	76.00 ± 0.67%
Transfer+GCN [3]	75.88 ± 0.23%	84.51 ± 0.13%
Transfer+Graph Interpolation	**76.47 ± 0.23%**	**85.23 ± 0.13%**

*GAT is evaluated with 600 test runs.

Table 3 - Evolution of the accuracy of few-shot classification on miniImageNet as a function of the number of supervised inputs s, and for various number of unsupervised queries q.

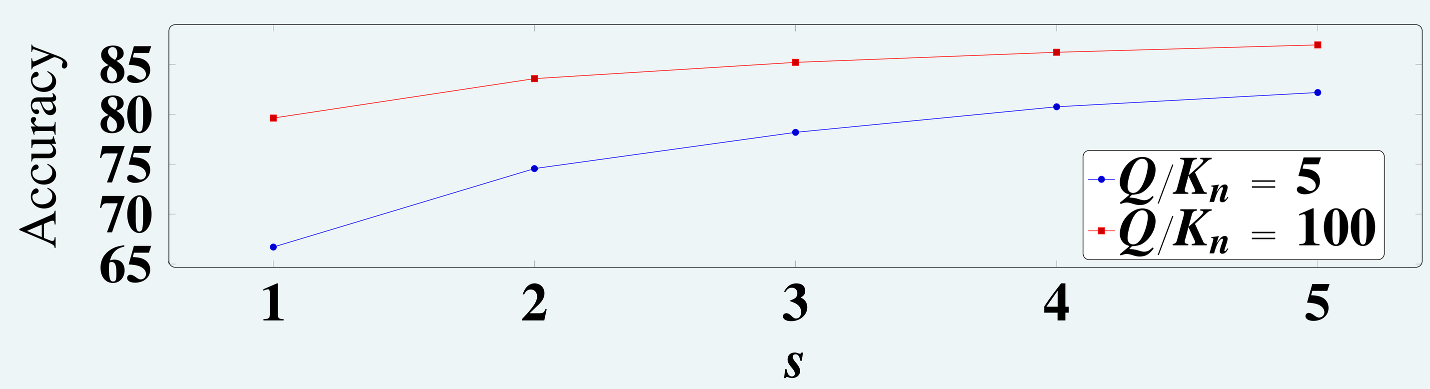

Conclusion

- We introduced a novel method to solve the few-shot classification problems. It consists in combining three steps : a pretrained backbone, a graph-based interpolation technique and a logistic regression.
- The proposed method obtain state-of-the-art results, with the most important gains in the case of 1-shot classification.
- The proposed method requires to tune few hyperparameters, and these have a little impact on accuracy. We thus believe that it is an applicable solution to many practical problems.
- There are still open questions to be addressed, such as the case of imbalanced classes, or settings where prediction must be performed on streaming data, one input at a time.

[1] C. Finn, P. Abbeel, and S. Levine.
Model-agnostic meta-learning for fast adaptation of deep networks.
In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.

[2] S. X. Hu, P. G. Moreno, Y. Xiao, X. Shen, G. Obozinski, N. D. Lawrence, and A. Damianou.
Empirical bayes transductive meta-learning with synthetic gradients.
arXiv preprint arXiv :2004.12696, 2020.

[3] T. N. Kipf and M. Welling.
Semi-supervised classification with graph convolutional networks.
arXiv preprint arXiv :1609.02907, 2016.

[4] P. Mangla, N. Kumari, A. Sinha, M. Singh, B. Krishnamurthy, and V. N. Balasubramanian.
Charting the right manifold : Manifold mixup for few-shot learning.
In *The IEEE Winter Conference on Applications of Computer Vision*, pages 2218–2227, 2020.

[5] J. Snell, K. Swersky, and R. Zemel.
Prototypical networks for few-shot learning.
In *Advances in Neural Information Processing Systems*, pages 4077–4087, 2017.

[6] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio.
Graph attention networks.
arXiv preprint arXiv :1710.10903, 2017.

[7] I. M. Ziko, J. Dolz, E. Granger, and I. B. Ayed.
Laplacian regularized few-shot learning.
arXiv preprint arXiv :2006.15486, 2020.