

Leveraging Class Hierarchies with Metric-Guided Prototype Learning

Vivien Sainte Fare Garnot
vivien.sainte-fare-garnot@ign.fr

Loic Landrieu
loic.landrieu@ign.fr

LASTIG, ENSG, IGN
Univ Gustave Eiffel
F-94160 Saint-Mande, France

Abstract

In many classification tasks, the set of target classes can be organized into a hierarchy. This structure induces a semantic distance between classes, and can be summarized under the form of a *cost matrix*, which defines a finite metric on the class set. In this paper, we propose to model the hierarchical class structure by integrating this metric in the supervision of a *prototypical network*. Our method relies on jointly learning a feature-extracting network and a set of class prototypes whose relative arrangement in the embedding space follows an hierarchical metric. We show that this approach allows for a consistent improvement of the error rate weighted by the cost matrix when compared to traditional methods and other prototype-based strategies. Furthermore, when the induced metric contains insight on the data structure, our method improves the overall precision as well. Experiments on four different public datasets—from agricultural time series classification to depth image semantic segmentation—validate our approach.

1 Introduction

Most classification models focus on maximizing the prediction accuracy, regardless of the semantic nature of errors. This can lead to high performing models, but puzzling errors such as confusing tigers and sofas, and casts doubt on what a model actually understands of the required task and data distribution. Neural networks in particular have been criticized for their tendency to produce improbable yet confident errors, notably when under adversarial attacks [1]. Training deep models to produce not only produce fewer but also *better* errors can increase their trustworthiness, which is crucial for downstream applications such as autonomous driving or land use and land cover monitoring [3, 11].

In many classification problems, the target classes can be organized according to a tree-shaped hierarchical structure. Such a taxonomy can be generated by domain experts, or automatically inferred from class names using the WordNet graph [34] or from word embeddings [33]. A step towards more reliable and interpretable algorithms would be to explicitly model the difference of gravity between errors, as defined by a hierarchical nomenclature.

For a classification task over a set \mathcal{K} of K classes, the hierarchy of errors can be encapsulated by a cost matrix $D \in \mathbb{R}_+^{K \times K}$, defined such that the cost of predicting class k when the

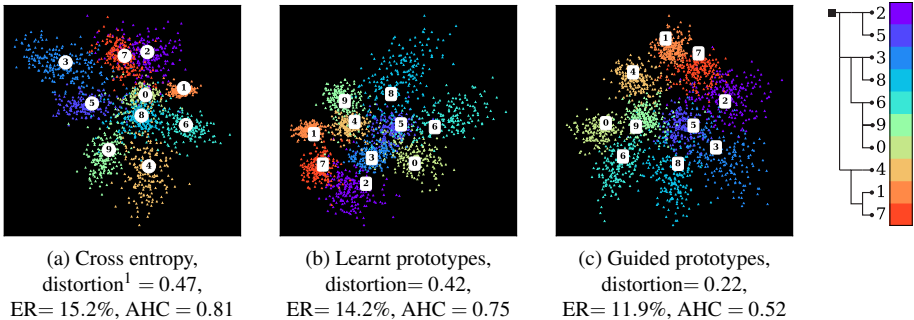


Figure 1: Mean class representation \circ , prototypes \square , and 2-dimensional embeddings \blacktriangle learnt on perturbed MNIST by a 3-layer convolutional net with three different classification modules: (a) cross-entropy, (b) learnt prototypes, and (c) learnt prototypes guided by a tree-shaped taxonomy (constructed according to the authors’ perceived visual similarity between digits). The guided prototypes (d) embed more faithfully the class hierarchy: classes with low error cost are closer. This is associated with a decrease in the *Average Hierarchical Cost* (AHC), as well as *Error Rate* (ER), indicating that our taxonomy may contain useful information for learning better visual features.

true class is l is $D[k, l] \geq 0$, and $D[k, k] = 0$ for all $k = 1 \cdots K$. Among many other options [27], one can define $D[k, l]$ as the length of the shortest path between the nodes corresponding to classes k and l in the tree-shaped class taxonomy.

As pointed out by Bertinetto et al. [3], the first step towards algorithms aware of hierarchical structures would be to generalize the use of cost-based metrics. For example, early iterations of the ImageNet challenge [11, 39] proposed to weight errors according to hierarchy-based costs. For a dataset indexed by \mathcal{N} , the *Average Hierarchical Cost* (AHC) between class predictions $y \in \mathcal{K}^{\mathcal{N}}$ and the true labels $z \in \mathcal{K}^{\mathcal{N}}$ is defined as:

$$\text{AHC}(y, z) = \frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} D[y_n, z_n]. \quad (1)$$

Along with the evaluation metrics, the loss functions should also take the cost matrix into account. While it is common to focus on retrieving certain classes through weighting [7, 29] or sampling [43, 47] schemes, preventing confusion between specific classes is less straightforward. For example, the cross entropy with one-hot target vectors singles out the predicted confidence for the true class, but treats all other classes equally. Beyond reducing the AHC, another advantage of incorporating the class hierarchy into the learning phase is that D may contain information about the structure of the data as well. Although it is not always the case, co-hyponyms (*i.e.* siblings) in a class hierarchy tend to share some structural properties. Encouraging such classes to have similar representations could lead to more efficient learning, *e.g.* by leveraging common feature detectors. Such priors on the class structure may be especially crucial when dealing with a large taxonomy, as noted by Deng et al. [11].

In this paper, we introduce a method to integrate a pre-defined class hierarchy into a classification algorithm. We propose a new distortion-based regularizer for prototypical network

¹For a formal definition of scale-free distortion, see Section 3.2; the distortion is computed with respect to the means of class embeddings for the cross entropy.

[9, 51]. This penalty allows the network to learn prototypes organized so that their pairwise distances reflect the error cost defined by a class hierarchy. Our contributions are as follows:

- We introduce a scale-independent formulation of the distortion between two metric spaces and an associated smooth regularizer.
- This formulation allows us to incorporate knowledge of the class hierarchy into a neural network at no extra cost in trainable parameters and computation.
- We show on four public datasets (CIFAR100, NYUDv2, S2-Agri, and iNaturalist-19) that our approach decreases the average cost of the prediction of standard backbones.
- As illustrated in Figure 1, we show that our approach can also lead to a better (un-weighted) precision, which we attribute to the useful priors contained in the hierarchy.

2 Related Work

Prototypical Networks: Our approach builds on the growing corpus of work on prototypical networks. These models are deep learning analogues of nearest centroid classifiers [48] and Learning Vector Quantization networks [26, 42], which associate to each class a representation, or prototype, and classify the observations according to the nearest prototype. These networks have been successfully used for few-shot learning [12, 45], zero-shot learning [23], and supervised classification [9, 16, 32, 51].

In most approaches, the prototypes are directly defined as the centroid of the learnt representations of samples of their classes, and updated at each episode [45] or iteration [16]. In the work of Mettes et al. [32] and Jetley et al. [23], the prototypes are defined prior to learning the embedding function. In this work, we follow the approach of [51] and learn the prototypes simultaneously with the data embedding function.

Hierarchical Priors: The idea of exploiting the latent taxonomic structure of semantic classes to improve the accuracy of a model has been extensively explored [44], from traditional Bayesian modeling [14, Chapter 5] to adaptive deep learning architectures [2, 38, 41, 50]. However, for these neural networks, the hierarchy is discovered by the network itself to improve the overall accuracy of the model. In our setting, the hierarchy is defined a priori and serves both to evaluate the quality of the model and to guide the learning process towards a reduced prediction cost.

Srivastava and Salakhutdinov [46] propose to implement Gaussian priors on the weight of neurons according to a fixed hierarchy. Redmon and Farhadi [37] implements an inference scheme based on a tree-shaped graphical model derived from a class taxonomy. Closest to our work, Hou et al. [21] propose a regularization based on the earth mover distance to penalize errors with high cost.

More recently, Bertinetto et al. [3] highlighted the relative lack of well-suited methods for dealing with hierarchical nomenclatures in the deep learning literature. They advocate for a more widespread use of the AHC for evaluating models, and detail two simple baseline classification modules able to decrease the AHC of deep models: *Soft-Labels* and *Hierarchical Cross-Entropy*. See the appendix for more details on these schemes. Following this objective, Karthik et al. [24] propose a an inference-time risk minimization scheme to reduce the AHC of the predictions based on the predicted posteriors.

Hyperbolic Prototypes: Motivated by their capacity to embed hierarchical data structures into low-dimensional spaces, [10], hyperbolic spaces are at the center of recent advances in modeling hierarchical relations [25, 36]. Closer to this work, [30, 31] also propose to embed a class hierarchy into the latent representation space. However, both approaches embed the class hierarchy before training the data embedding network. In contrast, we argue that incorporating the hierarchical structure during the training of the model allows the network and class embeddings to share their respective insights, leading to a better trade-off between AHC and accuracy. In this paper, we only explore Euclidean geometry, as this setting allows for the seamless integration of our method without changing the number of bits of precision or the optimizer [10].

Finite Metric Embeddings: Our objective of computing class representations with pairwise distances determined by a cost matrix has links with finding an isometric embedding of the cost matrix—seen as a finite metric. This problem has been extensively studied [4, 22] and is at the center of the growing interest for hyperbolic geometry [10]. Here, our goal is simply to influence the learning of prototypes with a metric rather than necessarily seeking the best possible isometry.

3 Method

We consider a generic dataset \mathcal{N} of N elements $x \in \mathcal{X}^{\mathcal{N}}$ with ground truth classes $z \in \mathcal{K}^{\mathcal{N}}$. The classes \mathcal{K} are organized along a tree-shape hierarchical structure, allowing us to define a cost matrix D by considering the shortest path between nodes. The matrix thus defined is symmetric, with a zero diagonal, strictly positive elsewhere, and respects the triangle inequality: $D[k, l] + D[l, m] \geq D[k, m]$ for all k, l, m in \mathcal{K} . In other words, D defines a finite metric. We denote by Ω an *embedding space* which, when equipped with the distance function $d : \Omega \times \Omega \mapsto \mathbb{R}_+$, forms a continuous metric space.

3.1 Prototypical Networks

A prototypical network is characterized by an embedding function $f : \mathcal{X} \mapsto \Omega$, typically a neural network, and a set $\pi \in \Omega^{\mathcal{K}}$ of K prototypes. π must be chosen such that any sample x_n of true class k has a representation $f(x_n)$ which is *close* to π_k and *far* from other prototypes.

Following the methodology of Snell et al. [45], a prototypical network (f, π) associates to an observation x_n the posterior probability over its class z_n defined as follows:

$$p(z_n = k | x_n) = \frac{\exp(-d(f(x_n), \pi_k))}{\sum_{l \in \mathcal{K}} \exp(-d(f(x_n), \pi_l))}, \forall k \in \mathcal{K} \quad (2)$$

We define an associated loss as the normalized negative log-likelihood of the true classes:

$$\mathcal{L}_{\text{data}}(f, \pi) = \frac{1}{N} \sum_{n \in \mathcal{N}} \left(d(f(x_n), \pi_{z_n}) + \log \left(\sum_{l \in \mathcal{K}} \exp(-d(f(x_n), \pi_l)) \right) \right). \quad (3)$$

This loss encourages the representation $f(x_n)$ to be close to the prototype of the class z_n and far from the other prototypes. Conversely, the prototype π_k is drawn towards the

representations $f(x_n)$ of samples n with true class k , and away from the representations of samples of other classes.

Following the insights of [51], the embedding function f and the prototypes π are learned simultaneously. This differs from many works on prototypical networks which learn prototypes separately or define them as centroids of representations. We take advantage of this joint training to learn prototypes which take into account both the distribution of the data and the relationships between classes, as described in the next section.

3.2 Metric-Guided Penalization

We propose to incorporate the cost matrix D into a regularization term in order to encourage the prototypes' positions in the embedding space Ω to be consistent with the finite metric defined by D . Since the sample representations are attracted to their respective prototypes in (3), such regularization will also affect the embedding network.

Metric Distortion As described in De Sa et al. [10], the distortion of a mapping $k \mapsto \pi_k$ between the finite metric space (\mathcal{K}, D) and the continuous metric space (Ω, d) can be defined as the average relative difference between distances in the source and target space:

$$\text{disto}(\pi, D) = \frac{1}{K(K-1)} \sum_{k, l \in \mathcal{K}^2, k \neq l} \frac{|d(\pi_k, \pi_l) - D[k, l]|}{D[k, l]}. \quad (4)$$

We argue that a network (f, π) trained to minimize $\mathcal{L}_{\text{data}}$ and whose prototypes π have a low distortion with respect to D should produce errors with low hierarchical costs. To understand the intuition behind this idea, let us consider a sample x_n of true class k and misclassified as class l . This tells us that the distance between $f(x_n)$ and π_l is small. If k and l have a high cost according to D , and since $k \mapsto \pi_k$ is of low distortion, then $d(\pi_k, \pi_l)$ must be large. The triangular inequality tells us that $d(f(x_n), \pi_k) \geq d(\pi_k, \pi_l) - d(f(x_n), \pi_l)$, and consequently that $d(f(x_n), \pi_k)$ must be large as well, which contradicts that (f, π) minimizes $\mathcal{L}_{\text{data}}$.

Scale-Free Distortion For a prototype arrangement π to have a small distortion with respect to a finite metric D as defined in Equation 4, the distance between prototypes must correspond to the distance between classes. This imposes a specific scale on the distances between prototypes in the embedding space. This scale may conflict with the second term of $\mathcal{L}_{\text{data}}$ which encourages the distance between embeddings and unrelated prototypes to be as large as possible. Therefore, lower distortion may also cause lower precision. To remove this conflicting incentive, we introduce a scale-independent formulation of the distortion (5) where $s \cdot \pi$ are the scaled prototypes, whose coordinates in Ω are multiplied by a scalar factor s . As shown in the appendix, $\text{disto}^{\text{scale-free}}$ can be efficiently computed algorithmically.

$$\text{disto}^{\text{scale-free}}(\pi, D) = \min_{s \in \mathbb{R}_+} \text{disto}(s \cdot \pi, D), \quad (5)$$

Distortion-Based Penalization We propose to incorporate the error qualification D into the prototypes' relative arrangement by encouraging a low *scale-free* distortion between π

and D . To this end, we define $\mathcal{L}_{\text{disto}}$, a smooth surrogate of $\text{disto}^{\text{scale-free}}$ (6), as detailed in the appendix, $\mathcal{L}_{\text{disto}}$ can be computed in closed form as a function of π and can be directly used as a regularizer.

$$\mathcal{L}_{\text{disto}}(\pi) = \frac{1}{K(K-1)} \min_{s \in \mathbb{R}_+} \sum_{k, l \in \mathcal{K}^2, k \neq l} \left(\frac{sd(\pi_k, \pi_l) - D[k, l]}{D[k, l]} \right)^2. \quad (6)$$

3.3 End-to-end Training

We combine $\mathcal{L}_{\text{data}}$ and $\mathcal{L}_{\text{disto}}$ in a single loss \mathcal{L} . $\mathcal{L}_{\text{data}}$ allows to jointly learn the embedding function f and the class prototypes π , while $\mathcal{L}_{\text{disto}}$ enforces a metric-consistent prototype arrangement, with $\lambda \in \mathbb{R}_+$ a hyper-parameter setting the strength of the regularization:

$$\mathcal{L}(f, \pi) = \mathcal{L}_{\text{data}}(f, \pi) + \lambda \mathcal{L}_{\text{disto}}(\pi). \quad (7)$$

4 Experiments

4.1 Datasets and Backbones

Table 1: Data composition and taxonomies of the four studied datasets. IR stands for the Imbalance Ratio (largest over smallest class count), nodes and leaves denote respectively the total number of classes and leaf-classes in the tree-shape hierarchy, ABF stands for the Average Branching Factor, and $\langle D \rangle$ stands for the average pairwise distance.

Dataset	Data			Hierarchical Tree			$\langle D \rangle$
	Volume (Gb)	Samples	IR	Depth	Nodes (leaves)	ABF	
NYUDv2	2.8	1449	93	3	57 (40)	5.0	4.3
S2-Agri	28.2	189 971	617	4	83 (45)	5.8	6.5
CIFAR100	0.2	60 000	1	5	134 (100)	3.8	7.0
iNat-19	82.0	265 213	31	7	1189 (1010)	6.6	11.0

We evaluate our approach with different tasks and public datasets with fine-grained class hierarchies: image classification on CIFAR100 [28] and iNaturalist-19 [49], RGB-D image segmentation on NYUDv2 [35], and image sequence classification on S2-Agri [40]. We define the cost matrix of these class sets as the length of the shortest path between nodes in the associated tree-shape taxonomies represented in the Appendix. As shown in Table 1, these datasets cover different settings in terms of data distribution and hierarchical structure.

Illustrative Example on MNIST: In Figure 1, we illustrate the difference in performance and embedding organization of the embedding space for different approaches. We use a small 3-layer convolutional net trained on MNIST with random rotations (up to 40 degrees) and affine transformations (up to 1.3 scaling). For plotting convenience, we set the features’ dimension to 2.

Image Classification on CIFAR100: We use a super-class system inspired by Krizhevsky et al. [28] and form a 5-level hierarchical nomenclature of size: 2, 4, 8, 20, and 100 classes. We use as backbone the established ResNet-18 [20] as embedding network for this dataset.

RGB-D Semantic Segmentation on NYUDv2: We use the standard split of 795 training and 654 testing pairs. We combine the 4 and 40 class nomenclatures of Gupta et al. [17] and the 13 class system defined by Handa et al. [18] to construct a 3-level hierarchy. We use FuseNet [19] as backbone for this dataset.

Image Sequence Classification on S2-Agri: S2-Agri is composed of 189 971 sequences of multi-spectral satellite images of agricultural parcels. We define a 4-level crop type hierarchy of size 4, 12, 19, and 44 classes with the help of experts from a European agricultural monitoring agency (ASP). We use the PSE+TAE architecture [40] as backbone, and follow their 5-fold cross-validation scheme for training. Crop mapping in particular benefits from predictions with a low hierarchical cost. Indeed, payment agencies monitor the allocation of agricultural subsidies and whether crop rotations follow best practice recommendations [15]. The monetary and environmental impact of misclassifications are typically reflected in the class hierarchy designed by domain experts [5, 6]. By achieving a low AHC, we ensure that these downstream tasks can be meaningfully realized from the predictions.

Fine-Grained Image Classification on iNaturalist-19 (iNat-19) iNat-19 [49] contains 1010 different classes organized into a 7 level hierarchy with respective width 3, 4, 9, 34, 57, 72, and 1010. We use ResNet-18 pretrained on ImageNet as backbone. We sample 75% of available images for training, while the rest is evenly split into a validation and test set.

4.2 Hyper-Parameterization

The embedding space Ω is chosen as \mathbb{R}^{512} for iNat-19 and \mathbb{R}^{64} for all other datasets. We chose d as the Euclidean norm. (see 4.6 for a discussion on this choice). We evaluate our approach (Guided-proto) with $\lambda = 1$ in (7) for all datasets. We use the same training schedules and learning rates as the backbone networks in their respective papers. In particular, the class imbalance of S2-Agri is handled with a focal loss [29].

4.3 Competing methods

In the paper where they are introduced, all backbone networks presented in Section 4.1 use a linear mapping between the sample representation and the class scores, as well as the cross-entropy loss. The resulting performance defines a baseline, denoted as `Cross-Entropy`, and is used to estimate the gains in Average Hierarchical Cost (AHC) and Error Rate (ER) provided by different approaches. We reimplemented other competing methods: Hierarchical Cross-Entropy (HXE) [3], Soft Labels [3], Earth Mover Distance regularization (XE+EMD) [21], Hierarchical Inference (YOLO) [37], Hyperspherical Prototypes (Hyperspherical-proto) [32], and Deep Mean Classifiers (Deep-NCM) [16]. See the Appendix for more details on these methods. Lastly, we evaluate simple prototype learning (Learnt-proto) [51] by setting $\lambda = 0$ in (7).

4.4 Analysis

Overall Performance: As displayed in Figure 2, the benefits provided by our approach can be appreciated on all datasets. Compared to the `Cross-Entropy` baseline, our model improves the AHC by 3% on NYUDv2 and S2-Agri, and up to 9% and 14% for CIFAR100, and

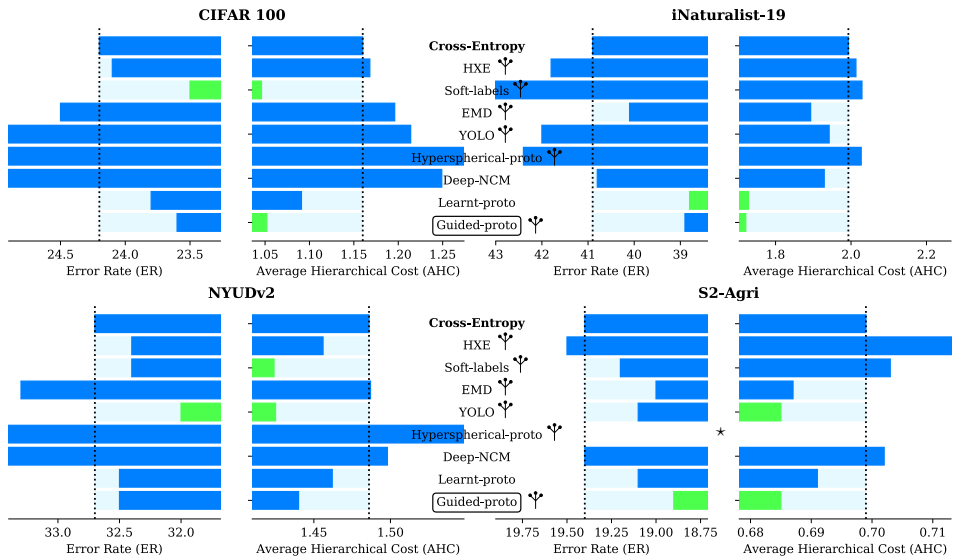


Figure 2: Error Rate (ER) in % and Average Hierarchical Cost (AHC) on four datasets for Guided-prototo, the Cross-Entropy baseline (in bold), and competing approaches. Methods that use the hierarchical knowledge are indicated with the symbol †. The best performances on each dataset are plotted in green. Our guided prototype approach improves both the ER and AHC across the four datasets compared to the baseline. The metrics are computed with the median over 5 runs for CIFAR100, the average over 5 cross-validation folds for S2-Agri, and a single run for NYUDv2 and iNat-19. The numeric values are given in the Appendix. (*: not evaluated).

iNat-19 respectively. The hierarchical inference scheme YOLO of Redmon and Farhadi [37] performs on par or better than our methods for NYUDv2 and S2-Agri, while Soft-labels perform well on CIFAR100 and NYUDv2. Yet, metric guided prototypes brings the most consistent reduction of the hierarchical cost across all tasks, datasets, and class hierarchies configurations. This suggest that arranging the embedding space consistently with the cost metric is a robust way of reducing a model’s hierarchical error cost. We argue that these results, combined with its ease of implementation, make a strong case for our approach.

While being initially designed to reduce the AHC, our method also provides a relative decrease of the ER by 3 to 4% across all datasets compared to the cross-entropy baseline. This indicates that cost matrices derived from the class hierarchies can indeed help neural networks to learn richer representations.

Prototype Learning: We observe that the learnt prototype approach Learnt-prototo consistently outperforms the Deep-NCM method. This suggests that defining prototypes as the centroids of their class representations might actually be disadvantageous. As illustrated on Figure 1, the positions of the embeddings tend to follow a Voronoi partition [13] with respect to the learnt prototypes of their true class rather than prototypes being the centroid of their associated representations. A surprising observation for us is that Learnt-prototo consistently outperforms the Cross-Entropy baseline, both in terms of AHC and ER.

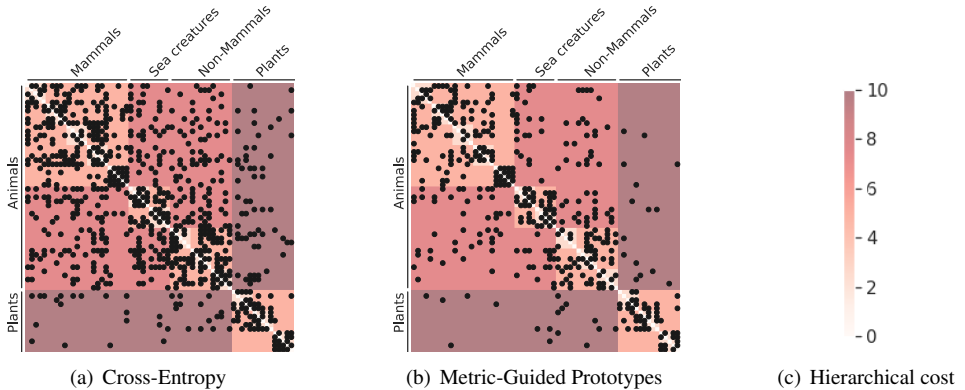


Figure 3: Partial confusion matrix for the “living organism” class subset of CIFAR100 for the Cross-Entropy baseline (a) and our approach (b). For readability, we only display (in black) entries of the matrices with at least one confusion. We also represent the cost of confusing different classes in shades of reds (c). We note that our approach yields fewer confusions between pairs of classes with high costs, such as plants and animals.

Computational Efficiency: Computing distances between representations and prototypes is comparable in terms of complexity than computing a linear mapping. The scaling factor in $\mathcal{L}_{\text{disto}}$ can be efficiently obtained as described in the Appendix. In practice, we observed that both training and inference time are identical for Cross-Entropy and Guided-PROTO: most of the time is taken by the computation of the embeddings.

4.5 Restricted Training Data Regime

We observed that the `Learnt-PROTO` method decreases the AHC across all four datasets even though it does not take the cost matrix into account. This suggests that, given enough data, this simple model can learn an empirical taxonomy through its prototypes’ arrangement. Furthermore, this taxonomy can share enough similarity with the one designed by experts to result in a decrease in AHC. To further evaluate the benefit of explicitly using the expert taxonomy with our approach, we train the models `Learnt-PROTO`, `Guided-PROTO`, and `EMD` with only part of the 160k images in the training set of `iNat-19`, and without pretraining on `ImageNet`. To compensate for the lack of data, we increase the regularization strength to $\lambda = 20$.




Figure 4: AHC of ResNet-18 trained on restricted training sets of `iNaturalist-19` with `Guided-PROTO`, `Learnt-PROTO`, and `EMD`. We represent the relative improvement compared to the performance of the `Cross-Entropy` baseline, which is shown on top of the plots.

In Figure 4, we observe that the two prototype-based approaches consistently improve the performance of the baseline for all training set sizes in terms of AHC. Moreover, the advantages brought by our proposed regularization are all the more significant when applied

to small training sets. This observation reinforces the idea that the `learnt-proto` method requires large amounts of data to learn a meaningful class hierarchy in an unsupervised way.

4.6 Ablation Study

In the Appendix, we present an extensive ablation study. We present here its main take-away.

Scale-Free Distortion: Our method for automatically choosing the best scale in our smooth distortion surrogate leads to an improvement of 0.9 ER on the iNat-19 dataset, which amounts to half the improvement compared to the baseline. In the other datasets, the improvements were more limited. We attribute the impact of our scale-free distortion on iNat-19 in particular to the structure of its class hierarchy: at the lowest level, iNat-19 classes have on average 14 co-hyponyms (siblings), compared to only 2 to 5 for the other datasets. When minimizing the distortion with a fixed scale of 1, the prototypes of hyponyms are incentivized to be close with respect to d since hyponyms have a small hierarchical distance of 2. This clashes with the minimization of the second part of $\mathcal{L}_{\text{data}}$ as defined in (3), which mutually repels prototypes of different classes. This conflict, made worse by classes with many hyponyms, is removed by our scale-free distortion. See the Appendix for additional insights into how our automatic scaling addresses this issue.

Choice of Metric Space: Prototypical networks operating on $\Omega = \mathbb{R}^m$ typically use the squared Euclidean norm in the distance function, motivated by its quality as a Bregman divergence [45]. However, given the large distance between prototypes induced by our regularization, this metric can cause stability issues. We observe for all datasets that that defining d as the Euclidean norm yields significantly better results across all datasets.

Guided vs. fixed prototypes : As suggested by the lower performance of `Hyperspherical-proto`, jointly learning the prototypes and the embedding network can be advantageous. To confirm this observation, we altered our `Guided-proto` method to first learn the prototypes and then the embedding network. We observed a significant decrease in performance across the board, up to 5 more points of ER in iNat-19. Conversely, we altered `Hyperspherical-proto` to learn spherical prototypes together with the embedding network. This improved the performance of `Hyperspherical-proto` even though it remained worse than the `Cross-Entropy` baseline (+1.20 ER, +0.10 AHC on CIFAR100). These observations suggest that insights from the data distribution can benefit the positioning of prototypes, and that they should be learned conjointly.

5 Conclusion

We introduced a new regularizer modeling the hierarchical relationships between the classes of a nomenclature. This approach can be incorporated into any classification network at no computational cost and with very little added code. We showed that our method consistently decreased the average hierarchical cost of three different backbone networks on different tasks and four datasets. Furthermore, our approach can reduce the rate of errors as well. In contrast to most recent works on hierarchical classification, we showed that this joint training is beneficial compared to the staged strategy of first positioning the prototypes and then training a feature extracting network. A PyTorch implementation of our framework as well as an illustrative notebook are available at <https://github.com/VSainteu/metric-guided-prototypes-pytorch>.

References

- [1] Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access*, 2018.
- [2] Shahanaz Ayub and JP Saini. ECG classification and abnormality detection using cascade forward neural network. *International Journal of Engineering, Science and Technology*, 2011.
- [3] Luca Bertinetto, Romain Mueller, Konstantinos Tertikas, Sina Samangooei, and Nicholas A Lord. Making better mistakes: Leveraging class hierarchies with deep networks. In *CVPR*, 2020.
- [4] Jean Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 1985.
- [5] Gerhard Brankatschk and Matthias Finkbeiner. Modeling crop rotation in agricultural lcas—challenges and potential solutions. *Agricultural Systems*, 2015.
- [6] Donald G Bullock. Crop rotation. *Critical reviews in plant sciences*, 1992.
- [7] Samuel Rota Bulo, Gerhard Neuhold, and Peter Kotschieder. Loss max-pooling for semantic image segmentation. In *CVPR*, 2017.
- [8] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *ICML*, 2005.
- [9] Chaofan Chen, Oscar Li, Daniel Tao, Alina Barnett, Cynthia Rudin, and Jonathan K Su. This looks like that: deep learning for interpretable image recognition. In *NeurIPS*, 2019.
- [10] Christopher De Sa, Albert Gu, Christopher Ré, and Frederic Sala. Representation trade-offs for hyperbolic embeddings. In *Proceedings of Machine Learning Research*, 2018.
- [11] Jia Deng, Alexander C Berg, Kai Li, and Li Fei-Fei. What does classifying more than 10,000 image categories tell us? In *ECCV*, 2010.
- [12] Nanqing Dong and Eric Xing. Few-shot semantic segmentation with prototype learning. In *BMVC*, 2018.
- [13] Steven Fortune. Voronoi diagrams and Delaunay triangulations. In *Computing in Euclidean Geometry*. World Scientific, 1992.
- [14] Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian data analysis*. CRC press, 2013.
- [15] Wyn Grant. *The common agricultural policy*. Macmillan International Higher Education, 1997.
- [16] Samantha Guerriero, Barbara Caputo, and Thomas Mensink. DeepNCM: deep nearest class mean classifiers. In *ICLR, Worskhop*, 2018.
- [17] Saurabh Gupta, Pablo Arbelaez, and Jitendra Malik. Perceptual organization and recognition of indoor scenes from RGB-D images. In *CVPR*, 2013.

- [18] A Handa, V Patraucean, V Badrinarayanan, S Stent, and R Cipolla. Understanding real world indoor scenes with synthetic data. In *CVPR*, 2016.
- [19] Caner Hazirbas, Lingni Ma, Csaba Domokos, and Daniel Cremers. Fusetnet: Incorporating depth into semantic segmentation via fusion-based CNN architecture. In *ACCV*, 2016.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [21] Le Hou, Chen-Ping Yu, and Dimitris Samaras. Squared earth mover’s distance-based loss for training deep neural networks. In *NeurIPS Workshop*, 2016.
- [22] Piotr Indyk, Jiří Matoušek, and Anastasios Sidiropoulos. Low-distortion embeddings of finite metric spaces. In *Handbook of discrete and computational geometry*. Chapman and Hall/CRC, 2017.
- [23] Saumya Jetley, Bernardino Romera-Paredes, Sadeep Jayasumana, and Philip Torr. Prototypical priors: From improving classification to zero-shot learning. In *BMVC*, 2015.
- [24] Shyamgopal Karthik, Ameya Prabhu, Puneet K Dokania, and Vineet Gandhi. No cost likelihood manipulation at test time for making better mistakes in deep networks. *ICLR*, 2021.
- [25] Valentin Khruikov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic image embeddings. In *CVPR*, 2020.
- [26] Teuvo Kohonen. Learning vector quantization. In *Self-organizing maps*. Springer, 1995.
- [27] Aris Kosmopoulos, Ioannis Partalas, Eric Gaussier, Georgios Paliouras, and Ion Androutsopoulos. Evaluation measures for hierarchical classification: a unified view and novel approaches. *Data Mining and Knowledge Discovery*, 2015.
- [28] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [29] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [30] Shaoteng Liu, Jingjing Chen, Liangming Pan, Chong-Wah Ngo, Tat-Seng Chua, and Yu-Gang Jiang. Hyperbolic visual embedding learning for zero-shot recognition. In *CVPR*, 2020.
- [31] Teng Long, Pascal Mettes, Heng Tao Shen, and Cees GM Snoek. Searching for actions on the hyperbole. In *CVPR*, 2020.
- [32] Pascal Mettes, Elise van der Pol, and Cees Snoek. Hyperspherical prototype networks. In *NeurIPS*, 2019.
- [33] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *ICLR Workshop*, 2013.

- [34] George A Miller, Richard Beckwith, Christiane Fellbaum, Derek Gross, and Katherine J Miller. Introduction to wordnet: An on-line lexical database. *International Journal of Lexicography*, 1990.
- [35] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from RGBD images. In *ECCV*, 2012.
- [36] Maximillian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *NeurIPS*, 2017.
- [37] Joseph Redmon and Ali Farhadi. YOLO9000: better, faster, stronger. In *CVPR*, 2017.
- [38] Deboleena Roy, Priyadarshini Panda, and Kaushik Roy. Tree-CNN: a hierarchical deep convolutional neural network for incremental learning. *Neural Networks*, 2020.
- [39] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 2015.
- [40] Vivien Sainte Fare Garnot, Loic Landrieu, Sebastien Giordano, and Nesrine Chehata. Satellite image time series classification with pixel-set encoders and temporal self-attention. In *CVPR*, 2020.
- [41] Ruslan Salakhutdinov, Joshua B Tenenbaum, and Antonio Torralba. Learning with hierarchical-deep models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2012.
- [42] Atsushi Sato and Keiji Yamada. Generalized learning vector quantization. *NeurIPS*, 1995.
- [43] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *CVPR*, 2016.
- [44] Carlos N Silla and Alex A Freitas. A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery*, 2011.
- [45] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *NeurIPS*, 2017.
- [46] Nitish Srivastava and Russ R Salakhutdinov. Discriminative transfer learning with tree-based priors. In *NeurIPS*, 2013.
- [47] Kah-Kay Sung. Learning and example selection for object and pattern detection. *MIT A.I.*, 1996.
- [48] Robert Tibshirani, Trevor Hastie, Balasubramanian Narasimhan, and Gilbert Chu. Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences*, 2002.
- [49] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The iNaturalist species classification and detection dataset. In *CVPR*, 2018.

- [50] Zhicheng Yan, Hao Zhang, Robinson Piramuthu, Vignesh Jagadeesh, Dennis DeCoste, Wei Di, and Yizhou Yu. HD-CNN: hierarchical deep convolutional neural networks for large scale visual recognition. In *ICCV*, 2015.
- [51] Hong-Ming Yang, Xu-Yao Zhang, Fei Yin, and Cheng-Lin Liu. Robust classification with convolutional prototype learning. In *CVPR*, 2018.